University of Alberta

REGIME DEPENDENT CONDITIONAL VOLATILITY IN FINANCIAL MARKETS

by

Lawrence L. Bauer ©

A thesis submitted to the Faculty of Graduate Studies and Research in partial fulfillment of the requirements for the degree of Doctor of Philosophy

in

Finance

Faculty of Business

Edmonton, Alberta
Fall 1999

National Library of Canada
Acquisitions and Bibliographic Services
395 Wellington Street
Ottawa ON K1A 0N4
Canada

Bibliothèque nationale du Canada
Acquisitions et services bibliographiques
395, rue Wellington
Ottawa ON K1A 0N4
Canada

Your file Votre référence

Our file Notre référence

The author has granted a non-exclusive licence allowing the National Library of Canada to reproduce, loan, distribute or sell copies of this thesis in microform, paper or electronic formats.

L'auteur a accordé une licence non exclusive permettant à la Bibliothèque nationale du Canada de reproduire, prêter, distribuer ou vendre des copies de cette thèse sous la forme de microfiche/film, de reproduction sur papier ou sur format électronique.

The author retains ownership of the copyright in this thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without the author's permission.

L'auteur conserve la propriété du droit d'auteur qui protège cette thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

0-612-46805-4

Canada

University of Alberta

Library Release Form

Name of Author: Lawrence L. Bauer

Title of Thesis: Regime Dependent Conditional Volatility in Financial Markets

Degree: Doctor of Philosophy

Year this Degree Granted: 1999

Permission is hereby granted to the University of Alberta Library to reproduce single copies of this thesis and to lend or sell such copies for private, scholarly, or scientific research purposes only.

The author reserves all other publication and other rights in association with the copyright in the thesis, and except as hereinbefore provided, neither the thesis nor any substantial portion thereof may be printed or otherwise reproduced in any material form whatever without the author's prior written permission.

37 Downing Street
St. John's, NF A1B 2R9
Canada

SEPTEMBER 28, 1999

Date submitted

University of Alberta

Faculty of Graduate Studies and Research

The undersigned certify that they have read, and recommend to the Faculty of Graduate Studies and Research for acceptance, a thesis entitled *Regime Dependent Conditional Volatility in Financial Markets* submitted by Lawrence L. Bauer in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Finance.

Dr. Steve Beveridge
Supervisor

Dr. Bob Korkie

Dr. Aditya Kaul

Dr. James Unterschultz

JD Jobson

Dr. David Jobson

I G Morgan

Dr. Ieuan Morgan
External Examiner

SEPTEMBER 27, 1999

Abstract

This study develops the Regime Dependent Generalized Auto Regressive Conditional Heteroskedasticity (RD-GARCH) model and applies it to a daily index of returns on U.S. equities covering the period 1926 to 1997. The model is based on the ARCH methodology first proposed by Engle (1982) and refined by subsequent researchers. The RD-GARCH model is different from previous models in that it combines the ARCH methodology with a general approach which allows model parameters to vary across periods or regimes of differing volatility.

Because traditional optimization techniques perform poorly when applied to highly non-linear models, including models from the GARCH family, a hybrid technique is developed which combines a stochastic global optimization algorithm with a deterministic local optimization algorithm. The hybrid technique exhibits greater robustness than the deterministic algorithm and converges to a solution more quickly than the stochastic algorithm.

Likelihood ratio tests of the within-sample properties of the RD-GARCH methodology demonstrate superiority of the model over a variety of models from the conventional GARCH family, both for the full sample period and for a number of sub-periods. The RD-GARCH model provides significantly improved model fit and also provides residuals with superior distributional properties.

The out-of-sample performance of the RD-GARCH methodology is compared to the performance of a number of alternative models for one-period ahead daily and monthly forecasts of variance. The RD-GARCH model is found to be the preferred model for forecasting both daily and monthly return variance.

Two alternative distributional assumptions for disturbances are investigated. The results for within-sample performance are inconclusive. The Student-t distribution is

found to provide improved out-of-sample forecasts for daily variance while the normal distribution is found to be the preferred distribution for monthly variance forecasts.

The results of this study indicate that combining a model which allows conditional volatility to change in an evolutionary manner with a model which allows for abrupt changes in unconditional volatility at discrete points in time provides a significant improvement over previously available models. The RD-GARCH model provides a general approach to combining these two alternative modeling methodologies, thereby providing a valuable addition to the GARCH family of conditional volatility models.

Dedication

To Shelley, whose support from the beginning made it possible.
To Katie, whose energy and joy of life made it bearable.
To Sam, whose appearance made it worthwhile.

To Joseph, whose presence and absence will always be felt.

Acknowledgements

This dissertation is the culmination of nearly four years direct effort, a further three years preparatory work in the Ph.D. program at the University of Alberta, and a lifetime of influences and experiences. It would not have been possible without the support, patience, and occasional kick in the pants from many people.

I must first thank my wife, Shelley, and our children, Katie and Sam, for their love, understanding, and support throughout the time I have spent in the Ph.D. program. I hope that I have not been an absent husband and father too often.

Given a finite amount of time and space, it is not possible to acknowledge everyone who has helped bring me to this completion of one part of my life and the threshold of a new part. There are a number of people, however, that I would like to mention individually.

I would like to begin by thanking my mother and father, June and Leonard Bauer, for instilling in me the curiosity to begin the Ph.D. program and the determination to complete it. I must also thank Boyd Bauerfind, a high-school electronics teacher who taught me a great deal about electronics and computers, but more about life. It is in no small part due to his example that I was motivated to seek out a profession that combines a quest for knowledge with the opportunity to share that knowledge.

It would clearly not have been possible to complete this endeavour without the suggestions and assistance provided by the members of my committee: Dave Jobson, Aditya Kaul, Bob Korkie, and Jim Unterschultz. I owe a special debt of gratitude to my supervisor, Steve Beveridge; his frequent ideas, suggestions, questions, support when necessary, and prodding when appropriate helped take the germ of an idea and turn it into this completed dissertation. I am also indebted to many of the students in the Ph.D. program for both moral and intellectual support over the past seven years. In particular, I would like to thank Ranjini Sivakumar and Angelo Unite, with whom I went through the economics and finance comprehensive exams, and whose support while studying made all the difference. I would also like to thank Jeanette Gosine, Administrator for the Ph.D. program at the University of Alberta, for her assistance with the myriad of administrative details associated with the program. Without her kind attention, my frequent late registrations, a program extension, and the postponement of my final defense due to the imminent birth of Samuel and Joseph would all have made my life much more difficult.

Finally, I would like to thank my father, Leonard Bauer, and Dr. Alex Faseruk of Memorial University of Newfoundland for proofreading the final manuscript of the dissertation and pointing out more editorial errors than should have remained at that stage. Any remaining errors are, of course, my own.

Contents

1 Introduction and Background 1
- 1.1 Introduction 1
- 1.2 Conditionally Heteroskedastic Models 2
  - 1.2.1 ARCH 3
  - 1.2.2 GARCH 3
  - 1.2.3 Other Members of the ARCH Family 4
    - Exponential GARCH 4
    - Nonlinear ARCH 5
    - APARCH 5
    - GJR-GARCH 6
    - N-A GARCH 6
    - AGARCH 7
  - 1.2.4 Nested GARCH 7
  - 1.2.5 Estimation of Nested GARCH Models 8
- 1.3 The News Impact Curve 9
- 1.4 Alternative Models 13
- 1.5 Regime Dependence 14

2 Estimation of GARCH Models 16
- 2.1 Optimization 16
  - 2.1.1 Deterministic vs. Stochastic Methods 17
- 2.2 Derivatives 21
  - 2.2.1 Analytic Derivatives 22
  - 2.2.2 Numeric Derivatives 22
  - 2.2.3 Automatic Differentiation 23

3 Model Selection 25
- 3.1 Replication of Previous Results 25
  - 3.1.1 Data 25
  - 3.1.2 Model 26
  - 3.1.3 Results 27
- 3.2 Model Selection 31
- 3.3 Residual Tests 34

4 Model Extensions **39**
4.1 Mean Equation Development 39
  4.1.1 Lagged Dependent Variables 40
  4.1.2 Weekend Effect 41
  4.1.3 Augmented Conditional Mean Equation 41
  4.1.4 Likelihood Equation 42
  4.1.5 Mean Equation Selection 42
4.2 Conditional Volatility Equation Development 44
  4.2.1 Regime Dependent GARCH 48
  4.2.2 Regime Detection 48
  4.2.3 Expanding Window Variance Ratio Regime Detection 50
  4.2.4 Correlation Among Explanatory Variables 52
4.3 The Final Model 52
4.4 Alternative Distributional Assumptions 53
  4.4.1 The Generalized Error Distribution 53
  4.4.2 The Student-*t* Distribution 54

5 Application of RD-GARCH **55**
5.1 Regime Detection 55
5.2 Behaviour of Returns During Regimes 59
5.3 Overall Performance of RD-GARCH 60
5.4 Further Analysis of the Selected Model 60
5.5 Tests of Alternative Distributions 68

6 Model Stability **75**
6.1 Return Series 75
  6.1.1 Stability of the Regime Selection Process 76
  6.1.2 Parameter Stability 78
6.2 Sub-sample Results 81
  6.2.1 Stability of the Regime Selection Process 82
  6.2.2 Parameter Stability 90
    Conventional Models 90
    Regime Dependent Models 101
    Summary 105

7 Out of Sample Forecasting **121**
7.1 A Note on Notation 121
7.2 Methodology 122
7.3 The Volatility of Returns 123
7.4 Forecasting Models 124
  7.4.1 Unconditional Variance 125
  7.4.2 Moving Average Variance 125
  7.4.3 Random Walk 126
  7.4.4 Conventional GARCH 127
  7.4.5 Extended GARCH 128

7.4.6 RD-GARCH 129
7.5 Forecast Evaluation 130
7.5.1 Loss Function Tests 131
Mean Error 131
Mean Absolute Error 131
Mean Square Error 132
7.5.2 Regression Tests 132
Test for Independent Information 132
Test for Difference in MSE 133
7.6 Results 135
7.6.1 Calibration Periods 135
7.6.2 Daily Results 137
Loss Function Tests 137
Tests for Sufficiency of RD-GARCH 140
Tests for Forecast Improvement 143
Tests for Difference in MSE 145
Summary 148
7.6.3 Monthly Results 149
Loss Function Tests 149
Tests for Sufficiency of RD-GARCH 153
Tests for Forecast Improvement 155
Tests for Difference in MSE 157
Summary 160
7.6.4 Alternative Distributional Assumptions 161

8 Conclusions **177**

Bibliography **184**

Appendices

A Model Summary **188**

B Regime Detection Example **191**

C Log-Likelihood Derivation **192**
C.1 Standard Normal Distribution 192
C.2 Generalized Error Distribution 193
C.3 Student-*t* Distribution 193

D Supplementary Tables **195**

List of Tables

3.1 Comparison of estimation techniques 27
3.2 Estimates of model parameters 28
3.3 LR tests for asymmetry 33
3.4 LR tests for model selection 34
3.5 Descriptive statistics for residuals 35

4.1 LR tests for alternative mean equations 43
4.2 LR tests for asymmetry (final mean equation) 44
4.3 Estimates of model parameters (final mean equation) 45
4.4 Tests of homogeneity of variance 47

5.1 Dates of regimes 56
5.2 Descriptive statistics for regimes 59
5.3 LR tests for RD-GARCH 61
5.4 Estimates of model parameters (final model) 62
5.5 Descriptive statistics for residuals (final model) 67
5.6 LR tests for alternative distributions 69
5.7 Estimates of model parameters (alternative distributions) 73
5.8 Descriptive statistics for residuals (alternative distributions) 74

6.1 Dates of regimes 1926–1990 (raw returns) 107
6.2 Dates of regimes 1926–1990 (excess returns) 108
6.3 Estimates of model parameters (raw returns, no regimes) 109
6.4 Estimates of model parameters (raw returns) 110
6.5 Dates of regimes 1915–1995 111
6.6 Dates of regimes 1915–1951 112
6.7 Dates of regimes 1952–1995 113
6.8 Dates of regimes 1915–1929 114
6.9 Dates of regimes 1929–1951 115
6.10 Dates of regimes 1952–1973 116
6.11 Dates of regimes 1974–1995 117
6.12 Estimates of model parameters (sub-sample, no regimes) 118
6.13 Estimates of model parameters (sub-sample) 119

7.1 Loss function tests (daily, 1952) 137
7.2 Loss function tests (daily, 1974) 139

7.3 Loss function tests (daily, rolling) 140
7.4 Tests for sufficiency of RD-GARCH (daily, 1952) 141
7.5 Tests for sufficiency of RD-GARCH (daily, 1974 and rolling) 142
7.6 Tests for forecast improvement (daily, 1952) 143
7.7 Tests for forecast improvement (daily, 1974 and rolling) 144
7.8 Tests for difference in MSE (daily, 1952) 146
7.9 Tests for difference in MSE (daily, 1974 and rolling) 147
7.10 Loss function tests (monthly, 1952) 150
7.11 Loss function tests (monthly, 1974) 151
7.12 Loss function tests (monthly, five year rolling) 152
7.13 Tests for sufficiency of RD-GARCH (monthly, 1952) 153
7.14 Tests for sufficiency of RD-GARCH (monthly, 1974 and rolling) . . . 154
7.15 Tests for forecast improvement (monthly, 1952) 155
7.16 Tests for forecast improvement (monthly, 1974 and rolling) 156
7.17 Tests for difference in MSE (monthly, 1952) 158
7.18 Tests for difference in MSE (monthly, 1974 and rolling) 159
7.19 Loss function tests (daily, alternative distributions) 162
7.20 Tests for difference in MSE (daily, alternative distributions) 162
7.21 Loss function tests (monthly, alternative distributions) 163
7.22 Tests for difference in MSE (monthly, alternative distributions) 163
7.23 Calibration parameters (GARCH, starting 1952) 165
7.24 Calibration parameters (extended GARCH, starting 1952) 166
7.25 Calibration parameters (RD-GARCH, starting 1952) 167
7.26 Calibration parameters (GARCH, starting 1974) 169
7.27 Calibration parameters (extended GARCH, starting 1974) 170
7.28 Calibration parameters (RD-GARCH, starting 1974) 171
7.29 Calibration parameters (GARCH, five year rolling) 173
7.30 Calibration parameters (extended GARCH, five year rolling) 174
7.31 Calibration parameters (RD-GARCH, five year rolling) 175

A.1 Summary of model definitions 189

D.1 Estimates of model parameters (alternative mean equations) 196
D.2 Descriptive statistics for residuals (alternative mean equations) 197
D.3 Estimates of model parameters (with regimes) 198
D.4 Descriptive statistics for residuals (with regimes) 203

List of Figures

1.1 The News Impact Curve: ω and α 9
1.2 The News Impact Curve: Asymmetry 10
1.3 The News Impact Curve: Shape 11
1.4 The History Impact Curve 12

4.1 Residuals from the conditional mean equation 46

5.1 Excess returns 1926–1990 57
5.2 News Impact Curve 1926–1990 64
5.3 History Impact Curve 1926–1990 66
5.4 News Impact Curve (alternative distributions) 70
5.5 History Impact Curve (alternative distributions) 71

6.1 Annualized monthly risk-free rate 1926–1990 75
6.2 Returns 1926–1990 77
6.3 News and History Impact Curves 1926–1990 (excess returns) 79
6.4 News and History Impact Curves 1926–1990 (raw returns) 80
6.5 Raw returns 1915–1995 84
6.6 Raw returns 1915–1951 87
6.7 Raw returns 1952–1995 88
6.8 News and History Impact Curves 1915–1995 93
6.9 News and History Impact Curves 1915–1951 94
6.10 News and History Impact Curves 1952–1995 95
6.11 News and History Impact Curves 1915–1929 97
6.12 News and History Impact Curves 1929–1951 98
6.13 News and History Impact Curves 1952–1973 99
6.14 News and History Impact Curves 1974–1995 100

Chapter 1

Introduction and Background

1.1 Introduction

Many techniques exist to model financial time series. When applied to a series comprising stock returns, these techniques provide estimates of the expected return or conditional mean of the series. Estimating the expected return of a financial series only provides part of the picture, as the usual assumption of risk aversion also requires specification of the variability of the expected return.

A number of techniques exist for specifying the nature of the expected volatility of financial series. Some of these techniques, such as the Auto Regressive Integrated Moving Average (ARIMA) approach of Box and Jenkins (1976), make the assumption that the variance of the underlying series is constant through time or that simple transformations such as taking logs or differencing will render the variance constant. Such assumptions are problematic, however, as inspection of stock return series over long time periods will indicate.

Alternatively, the model may be specified in such a way that the variance of the series being modeled is allowed to change over time. The number of available specifications for models with time varying volatility, or heteroskedasticity, is large, but a particularly useful family of such models has grown out of the Auto Regressive Conditional Heteroskedasticity (ARCH) specification first suggested by Engle (1982). Models from this family provide a specification that allows the conditional volatility of the series to evolve predictably through time as a function of current and prior unexpected shocks, prior conditional volatility, and possibly other exogenous variables. These models make the implicit assumption that the volatility "clumping" commonly observed in financial time series, such as stock returns, is the result of persistence in volatility from period to period; periods of high volatility tend to be followed by periods of high volatility, and periods of low volatility tend to be followed by periods of low volatility.

A third approach makes the assumption that the volatility undergoes abrupt shifts at discrete points in time. Such models (see for example Tsay, 1988) make the assumption that volatility in the market is based on exogenous factors which change abruptly. In essence, the models propose that volatility is dependent upon regimes in the market which are caused by factors outside of the financial series being modeled.

Both of these alternative explanations for time-varying volatility have merit, and some attempts have been made at combining the techniques to address specific modeling situations. This study seeks to combine the two techniques in a generalized manner to generate better estimates of expected volatility. In particular, the ARCH approach is combined with a technique that allows model parameters to depend upon the current exogenously determined unconditional volatility regime to form the Regime Dependent Generalized Auto Regressive Conditional Heteroskedasticity (RD-GARCH) model.

The remainder of this chapter is organized as follows: The next section introduces the ARCH family of conditionally heteroskedastic models and discusses a number of the alternative specifications which have been developed since Engle first suggested the methodology. Section 1.3 outlines the use of the News Impact Curve and introduces the History Impact Curve as tools for comparing alternative model specifications within the ARCH framework. Section 1.4 discusses two models which address the problem of time-varying variance in the context of shifts in the unconditional variance. The final section discusses several specialized models which combine ARCH effects with some form of regime dependence and outlines the organization of the remainder of this study.

1.2 Conditionally Heteroskedastic Models

The set of models of interest are all members of the broad class of heteroskedastic econometric models. These models consists of two parts: a mean equation which defines the evolution of a dependent variable in terms of one or more independent variables; and an equation which defines the nature of the time varying variance or heteroskedasticity. While the heteroskedasticity specification can take many forms, the particular set of models which are considered in this study treats heteroskedasticity as an evolutionary process. The model specification is two separate but related parts: The mean equation, and a conditional variance specification.

$$y_t = \gamma_0 + \sum_i \gamma_i x_i + u_t \tag{1.1}$$

The residual, u_t, has zero mean and variance h_t^2 which evolves through time:

$$h_t^2 = \omega + \sum_j \alpha_j z_j \tag{1.2}$$

Equation 1.1 is the mean equation and Equation 1.2 defines the conditional variance. In the above equations, y_t represents the series to be modeled, x_i represent the explanatory variables in the mean equation (which may, in addition to exogenous variables, include lagged values of the series and current or lagged values from the heteroskedasticity specification), and z_j represent the explanatory variables in the heteroskedasticity specification. Because the variance of u_t evolves through time, a useful decomposition of the residual is given by:

$$u_t = h_t v_t \tag{1.3}$$

or

$$v_t = \frac{u_t}{h_t} \tag{1.4}$$

where v_t is the standardized residual from the mean equation and has the useful property that $v_t \sim iid(0,1)$.

1.2.1 ARCH

The particular form of the conditional volatility to be considered is based upon Engle's (1982) ARCH model. The simple ARCH(q) model proposes that the variance of the residuals from the mean equation, u_t, evolves following the process:

$$\mathrm{var}\left(u_t|I_{t-1}\right) = h_t^2 = \omega + \sum_{i=1}^{q} \alpha_i u_{t-i}^2 \tag{1.5}$$

where $I_t = \{u_{t-s}, s \geq 0\}$. In this specification the conditional variance, h_t^2, is a function of q lagged squared residuals from the mean equation.

One problem with the ARCH specification is that large values of q are frequently required to achieve a satisfactory fit for a particular dataset. Such models generally require an artificial fixed lag structure which forces declining lag weights on the model parameters to avoid parameter estimates which cause negative values for the conditional variance (Engle, 1982).

1.2.2 GARCH

The Generalized Autoregressive Conditional Heteroskedasticity (GARCH) model of Bollerslev (1986) seeks to resolve this problem by including lagged values of the conditional variance directly in the conditional variance specification. This leads to the GARCH(p,q) specification for the conditional variance of u_t:

$$\mathrm{var}\left(u_t|I_{t-1}\right) = h_t^2 = \omega + \sum_{i=1}^{q} \alpha_i u_{t-i}^2 + \sum_{j=1}^{p} \beta_j h_{t-i}^2 \tag{1.6}$$

where again $I_t = \{u_{t-s}, s \geq 0\}$. Here, the conditional variance, h_t^2, is a function of q lagged squared residuals and p lagged conditional variance terms. This specification frequently requires fewer parameters, offers improved fit, and does not require the artificial parameter structure required by Engle's simple ARCH model. A form of the model which is particularly common is the GARCH(1,1) specification:

$$h_t^2 = \omega + \alpha u_{t-1}^2 + \beta h_{t-1}^2 \tag{1.7}$$

1.2.3 Other Members of the ARCH Family

While Bollerslev's GARCH specification provides one of the early extensions of the ARCH framework, it is not unique. Other extensions to the basic ARCH model include, the Exponential GARCH, Nonlinear ARCH, Asymmetric Power ARCH, GJR-GARCH, Nonlinear Asymmetric GARCH, and Asymmetric absolute value GARCH models. These models are discussed in more detail below.

Exponential GARCH

Nelson (1991) identifies and addresses two problems with GARCH models: non-negativity and asymmetry.

Non-negativity: It is clearly necessary with GARCH models that the conditional variance, h_t^2, be strictly non-negative. Many forms of the GARCH conditional volatility specification enforce this requirement via parameter restrictions.[1] The Exponential GARCH (EGARCH) model avoids the need for parameter restrictions by allowing the conditional volatility to evolve as the log of the conditional variance, $\ln h_t^2$.

Asymmetry: Both the conventional ARCH and the conventional GARCH models result in a conditional volatility specification that is a symmetric function of the residuals from the mean equation. The impact on the conditional variance of an unexpected shock depends only upon the magnitude of the shock, not its sign. The EGARCH specification allows the conditional variance to be an asymmetric function of the residuals from the mean equation—that is, positive and negative shocks of equal size can have a different impact on subsequent conditional volatility. The conditional volatility specification for the EGARCH model is:

$$\ln h_t^2 = \omega + \sum_{i=1}^{\infty} \alpha_i g(v_{t-i}) \tag{1.8}$$

where

$$g(v_t) \equiv |v_t| - E|v_t| - cv_t \tag{1.9}$$

$g(v_t)$ represents a linear combination of v_t and $|v_t|$ in which both v_t and $|v_t| - E|v_t|$ have zero mean. The EGARCH(1,1) form of Equation 1.8 can also be expressed in terms of finite lags as:

---

[1] Bollerslev's GARCH specification, for example, requires $\omega \geq 0$ in Equation 1.6

$$\ln h_t^2 = \omega + \alpha_1 \left[|v_{t-1}| - E|v_{t-1}| - c v_{t-1} \right] + \beta \ln h_{t-1}^2 \tag{1.10}$$

In the EGARCH specification, the parameter c controls the asymmetry in the response of the conditional volatility to unexpected shocks from the mean equation. A positive value for c will cause a negative shock (negative value for v_{t-1}) to have a larger impact on $\ln h_t^2$ than a positive shock of the same magnitude. Similarly, a negative value for c will cause a positive shock to have a larger impact than a negative shock of the same magnitude.

Nonlinear ARCH

The Nonlinear ARCH (NARCH) model of Higgins and Bera (1992) provides a more general functional form of the ARCH model which allows for nonlinear evolution of the conditional volatility:

$$h_t^2 = \left[\phi_0 (\sigma_u^2)^\delta + \sum_{i=1}^{q} \phi_i (u_{t-i}^2)^\delta \right]^{\frac{1}{\delta}} \tag{1.11}$$

where σ_u^2 is the unconditional variance of the residuals from the mean equation and

$$\begin{aligned} &\delta > 0 \\ &\phi_i \geq 0 \qquad \text{for} \quad i = 0, 1, \ldots, p \end{aligned}$$

in addition

$$\sum_{i=0}^{q} \phi_i = 1$$

Rearranging Equation 1.11 and combining terms results in the Box-Cox power transformation form of the model:

$$\frac{h_t^\lambda - 1}{\lambda} = \omega + \sum_{i=1}^{q} \alpha_i \frac{|u_{t-i}|^\lambda - 1}{\lambda} \tag{1.12}$$

APARCH

The Asymmetric Power ARCH (APARCH) model of Ding et al. (1993) incorporates the Box-Cox transformation of the conditional volatility and also allows for asymmetry in the response of conditional volatility to unexpected shocks. The functional form of this model is:

$$h_t^{\lambda} = \omega + \sum_{i=1}^{q} \alpha_i \left(|u_{t-i}| - c_i u_{t-i}\right)^{\lambda} + \sum_{j=1}^{p} \beta_j h_{t-j}^{\lambda} \tag{1.13}$$

The specification incorporates the power transformation properties of the NARCH model and also allows for asymmetry in the response of the conditional volatility to shocks from the mean equation in a manner similar to that provided in the EGARCH model.

GJR-GARCH

As with the EGARCH and APARCH models, the GARCH(1,1) model of Glosten et al. (1993) (GJR-GARCH) provides for asymmetry. In addition, the model allows the risk-free interest rate to impact the conditional volatility directly as an additional explanatory variable. The functional form of the model is:

$$h_t^2 = \omega + \gamma r_{ft} + \phi_1 u_{t-1}^2 + \phi_2 u_{t-1}^2 S_{t-1}^- + \beta h_{t-1}^2 \tag{1.14}$$

where r_{ft} represents the risk-free interest rate and S_{t-1}^- is a $(0, 1)$ dummy which is one when u_{t-1} is negative and zero otherwise. Manipulating Equation 1.14:

$$\begin{aligned} h_t^2 &= \omega + \gamma r_{ft} + \phi_1 u_{t-1}^2 + \phi_2 \left(u_{t-1}^2 - u_{t-1}|u_{t-1}|\right) + \beta h_{t-1}^2 \\ &= \omega + \gamma r_{ft} + u_{t-1} \left(\phi_1 u_{t-1} + \phi_2 u_{t-1} - \phi_2 |u_{t-1}|\right) + \beta h_{t-1}^2 \\ &= \omega + \gamma r_{ft} + \alpha_1 u_{t-1} \left(|u_{t-1}| - c u_{t-1}\right) + \beta h_{t-1}^2 \end{aligned} \tag{1.15}$$

demonstrates the similarity in the form of asymmetry in the GJR-GARCH model to that found in the EGARCH and APARCH specifications.

N-A GARCH

The Nonlinear Asymmetric GARCH (N-A GARCH) specification of Engle and Ng (1993) also allows the conditional variance to be an asymmetric function of the residuals from the mean equation. The conditional variance specification for the model is:

$$h_t^2 = \omega + \alpha_1 \left(u_{t-1} + b h_{t-1}\right)^2 + \beta h_{t-1}^2 \tag{1.16}$$

In this model, the parameter b controls the extent of the asymmetry in the response of the conditional volatility to shocks of equal magnitude but different sign. To illustrate the nature of the asymmetry, it is useful to rewrite Equation 1.16 as:

$$\begin{aligned} h_t^2 &= \omega + \alpha_1 \left(v_{t-1} h_{t-1} + b h_{t-1}\right)^2 + \beta h_{t-1}^2 \\ &= \omega + \alpha_1 h_{t-1}^2 \left(v_{t-1} + b\right)^2 + \beta h_{t-1}^2 \end{aligned} \tag{1.17}$$

As Equation 1.17 illustrates, a positive value for b will cause the impact on subsequent conditional volatility, h_t, of a positive shock (positive value of v_{t-1}) to be larger than the impact of a negative shock of the same magnitude. Similarly, a negative value for b will cause the impact on subsequent conditional volatility to be larger for a negative shock than for a positive shock of the same magnitude. The form of asymmetry provided in the N-A GARCH model is different than the form provided in the EGARCH, APARCH, and GJR-GARCH models, however. The nature of this difference is further addressed in Section 1.3.

AGARCH

The Asymmetric absolute value GARCH (AGARCH) model of Hentschel (1995) provides a specification for volatility which evolves through the square root of the conditional variance. In addition, the model also allows for asymmetry in the response of the conditional volatility to shocks of equal magnitude but different sign. The form of this model is:

$$h_t = \omega + \alpha_1 \left[|v_{t-1} - b| - c\,(v_{t-1} - b) \right] + \beta h_{t-1} \tag{1.18}$$

The model allows for asymmetry in the form provided in the EGARCH, APARCH, and GJR-GARCH models as well as the form provided in the N-A GARCH model.

There are many other members of this family of conditionally heteroskedastic models. Engle and Rothschild (1992) and Bollerslev et al. (1992) provide reviews of this group of models. In the interests of clarity, the term GARCH is used for the remainder of this study to refer to all members of the general class of ARCH models.

1.2.4 Nested GARCH

The proliferation of alternative and sometimes partially overlapping specifications for GARCH models leads to a problem in that it is often not possible to determine easily where one model leaves off and another starts. Hentschel (1995) provides a single specification which nests several of the more popular extensions to the GARCH model directly into a single equation. The form of this equation for the GARCH(1,1) family of models is:

$$\frac{h_t^{\lambda} - 1}{\lambda} = \omega' + \alpha' h_{t-1}^{\lambda} f^{\nu}(v_{t-1}) + \beta \frac{h_{t-1}^{\lambda} - 1}{\lambda} \tag{1.19}$$

where

$$f(v_t) = |v_t - b| - c(v_t - b) \tag{1.20}$$

and v_t represents the standardized residual obtained from the mean equation as defined in Equation 1.4 above.

The sub-models which are nested in this specification obtain from equality and inequality constraints on the parameters of Equations 1.19 and 1.20. For example, under the restriction $\lambda = 2$, $\nu = 2$, $b = 0$, and $c = 0$, Equation 1.20 reduces to:

$$f(v_t) = |v_t| \tag{1.21}$$

and Equation 1.19 reduces to:

$$\frac{h_t^2 - 1}{2} = \omega' + \alpha' h_{t-1}^2 v_{t-1}^2 + \beta \frac{h_{t-1}^2 - 1}{2} \tag{1.22}$$

or

$$\begin{aligned} h_t^2 &= 2\omega' - \beta + 1 + 2\alpha' h_{t-1}^2 v_{t-1}^2 + \beta h_{t-1}^2 \\ &= \omega + \alpha u_{t-1}^2 + \beta h_{t-1}^2 \end{aligned} \tag{1.23}$$

where

$$\begin{aligned} \omega &= 1 + 2\omega' - \beta \\ \alpha &= 2\alpha' \end{aligned}$$

which is Bollerslev's GARCH(1,1) specification. Other models from Section 1.2.3 can be similarly constructed by imposing parameter restrictions. The appropriateness of a particular nested model to represent a particular data series can then be assessed using a likelihood ratio (LR) test of the parameter restrictions.

1.2.5 Estimation of Nested GARCH Models

As the functional form of the Nested GARCH specification provided in Section 1.2.4 does not lend itself well to estimation, it is useful to rearrange the equation. For the case where $\lambda \neq 0$, Equation 1.19 can be written:

$$\begin{aligned} h_t &= \left(1 + \omega'\lambda - \beta + \alpha'\lambda h_{t-1}^{\lambda} f^{\nu}(v_{t-1}) + \beta h_{t-1}^{\lambda}\right)^{\frac{1}{\lambda}} \\ &= \left(\omega + \alpha h_{t-1}^{\lambda} f^{\nu}(v_{t-1}) + \beta h_{t-1}^{\lambda}\right)^{\frac{1}{\lambda}} \end{aligned} \tag{1.24}$$

where

$$\begin{aligned} f(v_t) &= |v_t - b| - c(v_t - b) \\ \omega &= 1 + \omega'\lambda - \beta \\ \alpha &= \alpha'\lambda \end{aligned}$$

For the case where $\lambda = 0$, using l'Hôpital's rule, Equation 1.19 can be written as:

$$h_t = \exp(\omega + \alpha f^{\nu}(v_{t-1}) + \beta \log h_{t-1})) \tag{1.25}$$

Figure 1.1: The News Impact Curve: ω and α

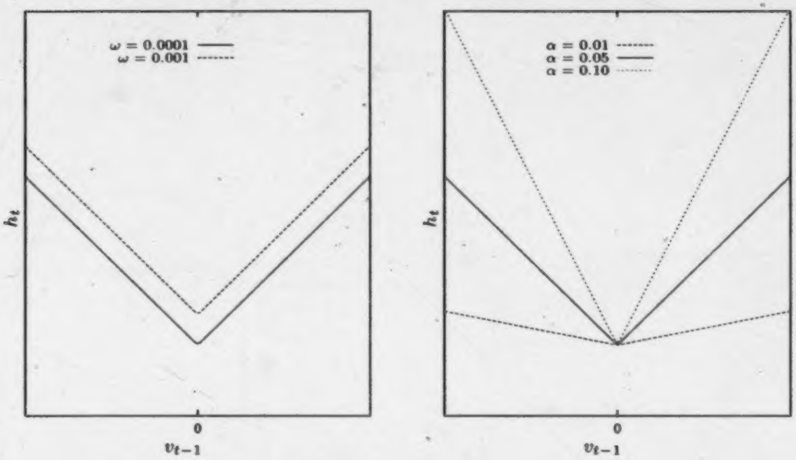

The figure shows example News Impact Curves which demonstrate the effect of changing the parameters ω and α. The parameter ω locates the NIC in space, and represents the unconditional volatility. The parameter α controls the curvature of the NIC and represents the short-term memory of conditional volatility.

1.3 The News Impact Curve

While the nested GARCH specification is attractive in that it allows simple likelihood ratio tests to determine the most appropriate model for a particular modeling situation, the large number of parameters provided can make the results difficult to interpret. One tool for simplifying the interpretation is the *News Impact Curve* (NIC) of Engle and Ng (1993). The NIC relates the conditional volatility, h_t, to the lagged standardized residuals of the mean equation, v_{t-1}, providing a graphical indication of the impact of unexpected shocks, or news,[2] on subsequent conditional volatility. The curve indicates how positive and negative shocks of differing magnitudes are impounded into conditional volatility. If the unexpected shocks are interpreted as news, the sign represents the type of news: a positive shock represents "good news" and a negative shock indicates "bad news". Similarly, the magnitude of the shock represents the size of the news event: a large negative shock indicates that the market has received a piece of very bad news, a smaller negative shock indicates that the market did not judge the news to be as important. The various parameters of the model affect this curve in different ways.

Figure 1.1 shows the effect of the parameters ω and α on the NIC. These are the parameters which appear in Engle's original ARCH conditional variance specification, and correspond to the intercept and the coefficient of the lagged residual from the mean equation (Equation 1.5 with $q = 1$). The intercept, ω, locates the NIC in space

---

[2] Throughout the discussion in this section and subsequent chapters, the terms "news" and "unexpected shocks" will be used interchangeably to refer to the residuals from the mean equation.

Figure 1.2: The News Impact Curve: Asymmetry

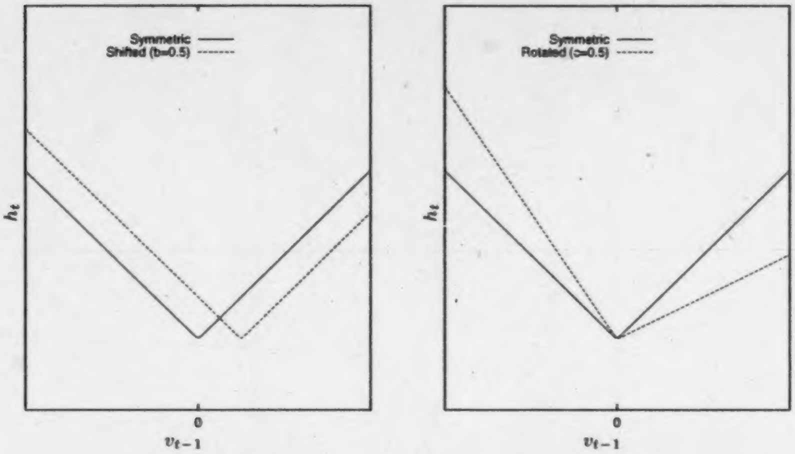

The figure shows example News Impact Curves which demonstrate the effect of changing the asymmetry parameters, b and c. The parameter b provides for asymmetry in the response of conditional volatility to unexpected shocks in the prior period by allowing the NIC to shift. The parameter c provides for asymmetry in the response of conditional volatility to unexpected shocks in the prior period by allowing for rotation of the NIC.

and represents unconditional volatility. The parameter α represents the short term memory of the conditional volatility specification by controlling the curvature of the NIC. Smaller values of α are represented by a flatter NIC, which indicates that the impact on conditional volatility in the following period for a given piece of news, or news event, is smaller for a model with a small value for α than for a model with a large value for α.

For the simple ARCH and GARCH specifications of Engle and Bollerslev respectively, the impact of news on volatility is symmetric. Equally large positive and negative shocks have the same impact on conditional volatility in the next period. As shown in Figure 1.2, different GARCH specifications allow for asymmetry in the response of conditional volatility to news in two ways: The NIC may be shifted, or it may be rotated.

The parameter b in the nested GARCH specification allows for the shifting of the NIC. This is the method provided in the N-A GARCH model. With $b > 0$ the NIC is shifted to the right (as shown). A negative shock of one standard deviation will have a larger impact on subsequent volatility than a positive shock of one standard deviation. That is to say, a piece of bad news will cause more subsequent uncertainty than a piece of good news of the same magnitude. While this approach to modeling asymmetry is later shown to perform well,[3] it is difficult to provide an economic interpretation for a shift in the NIC. A shift of the curve to the right by an amount b implies that no news ($v_t = 0$) results in greater subsequent conditional volatility than

---

[3] See Section 3.2, and in particular Tables 3.3 and 3.4.

Figure 1.3: The News Impact Curve: Shape

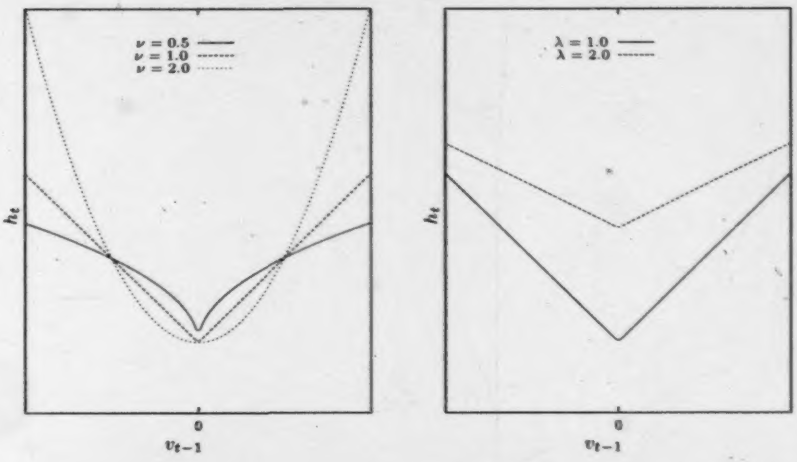

The figure shows example News Impact Curves which demonstrates the effect of changing the parameters ν and λ. The parameter ν controls the shape of the NIC, and thus the relationship between the impact of large and small shocks on subsequent conditional volatility. The parameter λ controls the Box-Cox transformation of the conditional volatility series, and does not have a simple, well-defined impact on the NIC. The impact appears to be similar to a combination of the impact of ω and α, which is to say that larger values for λ are characterized by a NIC which is located higher in space and is less curved, however this appearance is associated with a transformation of the conditional volatility series rather than a simple change in the nature of the impact of current shocks on subsequent conditional volatility.

some positive unexpected shock ($0 < v_t < b$).

The parameter c in the nested GARCH specification allows for the rotation of the NIC. This is the method provided in the EGARCH, APARCH, and GJR-GARCH models. In this case $c > 0$ indicates a clockwise rotation of the curve, which represents an effect that is somewhat similar to a shift of the curve to the right. Once again a negative shock will have a larger impact than an equally large positive shock. Because the curve is rotated rather than shifted, however, this effect is magnified for progressively larger shocks. The difference in impact for a negative shock with a magnitude of two standard deviations and a negative shock with a magnitude of one standard deviation will be larger than the difference in impact for a positive shock with a magnitude of two standard deviations and a positive shock with a magnitude of one standard deviation. Of course, these asymmetry effects can be combined—the NIC can be both shifted and rotated simultaneously.

As shown in Figure 1.3, the parameter ν controls the shape of the NIC. For the case where $\nu = 1$, the relationship between small and large shocks is linear. The difference in impact for shocks with a given difference in magnitude is the same, regardless of the magnitude of the shocks. When $\nu < 1$, the curve is concave and the difference in impact for shocks with a given difference in magnitude is greater for small shocks than for large shocks. For example, the difference in the impact

Figure 1.4: The History Impact Curve

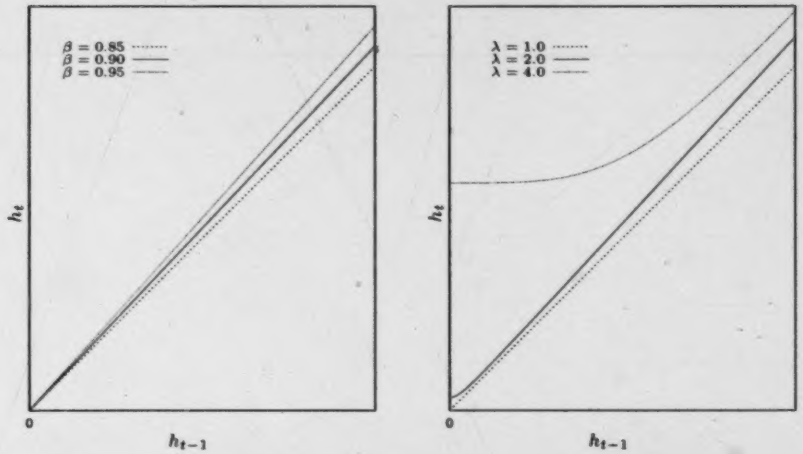

The figure shows example History Impact Curves which demonstrate the effect of changing the parameters β and λ. The parameter β controls the slope of the HIC, and has an effect similar to the parameter α in the NIC shown in the right-hand panel of Figure 1.1. The parameter λ controls the shape of the HIC in a manner similar to the effect of parameter ν on the NIC. As outlined in reference to the NIC, λ controls the transformation of the conditional volatility series itself. The effect noted on the shape of the HIC is only part its impact.

on subsequent volatility between a shock with a magnitude of $v_{t-1} = 1$ and a shock with a magnitude of $v_{t-1} = 2$ will be greater than the difference in the impact on subsequent volatility between a shock with a magnitude of $v_{t-1} = 2$ and a shock with a magnitude of $v_{t-1} = 3$. Similarly, when $\nu > 1$, the curve is convex and difference in impact for shocks with a given difference in magnitude is greater for large shocks than for small shocks.

The parameter β does not directly affect the NIC. Instead, it describes the impact of prior period conditional volatility on current conditional volatility. In essence, this parameter defines the long-term memory of the conditional volatility model. As it is not possible to interpret this effect in relation to the NIC, an interpretation is afforded by a new device called the *History Impact Curve* (HIC). It is so titled because it represents the contribution of the conditional volatility from the previous period to the current conditional volatility. Since the impact of the shock in the immediately prior period does not directly impact this curve, it represents longer term effects, or the history of the conditional volatility, rather than news events. Figure 1.4 shows that the effect of β in the HIC is similar to the effect of α in the NIC; it controls the slope of the curve. This parameter controls the degree of persistence in the conditional volatility. The greater the slope of the HIC, the more persistent the conditional volatility.

The parameter λ defines a Box-Cox transformation of the conditional volatility series. As such it can be interpreted in terms of specifying the functional form for the representation of the conditional volatility. Judge et al. (1985), pp. 839–842, suggest that the Box-Cox transformation may be useful in transforming a series which is non-normally distributed and heteroskedastic into one which is more nearly homoskedastic and normally distributed, but note that when λ is estimated as a parameter of the model:

> ... the estimated λ will be biased in the direction required for the transformed dependent variable to be more nearly homoskedastic ...

While the impact of λ cannot be defined simply in terms of how it affects the NIC or the HIC, its effect can readily be demonstrated in both of these devices. The right-hand panel of Figure 1.3 shows that the effect of λ on the NIC is similar to the combined effect of the parameters ω and α. Increasing values for λ are associated with a NIC which is located higher in space and exhibits less curvature. The right-hand panel of Figure 1.4 shows that λ affects the shape of the HIC in similar manner to the effect of ν in the NIC. The parameter λ allows for different levels of persistence for small and large values of conditional volatility.

1.4 Alternative Models

The broad class of GARCH models outlined previously represents one approach to the problem of modeling time varying volatility. These models consider the conditional volatility to be an evolutionary process in which the current conditional volatility is dependent upon prior conditional volatility and upon unexpected shocks in prior periods. Since the parameters of these models do not change through time, there is an implicit assumption that the unconditional volatility, σ_u^2, is constant.[4]

An alternative is to provide a model which allows the unconditional volatility to shift at discrete points. This methodology assumes that the series being modeled is generated by a process which shifts between periods or regimes characterized by differing unconditional volatility. One approach, used by Tsay (1988), fits a simple ARIMA model to the series and looks for an abrupt shift in the variance of the residuals. If such a shift is detected, a transformation is applied to the original series to remove the impact of the variance shift and the model is re-estimated. The process is repeated until there are no more variance shifts detected in the residuals from the ARIMA model. Hamilton (1989) provides a somewhat different approach through a Markov switching model applied to the series. This model generates estimates of the probabilities that the series has switched into one of several predefined volatility states.

---

[4] For the simple ARCH specification of Engle (1982) (Equation 1.5) the unconditional volatility of the residuals from the mean equation, $\sigma_u^2 = E(u_t^2) = E(h_t^2)$, is given by:

$$\sigma_u^2 = \frac{\omega}{1 - \sum_{i=1}^{q} \alpha_i} \tag{1.26}$$

If the α_i are time-invariant, the unconditional volatility will also be constant through time.

1.5 Regime Dependence

There have been attempts made at combining GARCH models with models which allow for shifts in the unconditional volatility of the series. In particular, Hamilton and Susmel (1994) combine the ARCH framework with Hamilton's Markov switching model. This approach uses the Markov switching model to determine the probability that the economy is in a particular volatility state (regime) and then allows some of the parameters of the ARCH conditional volatility equation to vary depending upon the regime. Hamilton and Susmel's approach only works with ARCH models, as the lagged conditional volatility from the GARCH model requires the Markov chain to be maintained for the total number of observations over which the model is estimated, rendering the approach numerically intractable for GARCH models covering a large number of observations.

Gray (1996) extends the Markov switching process to include a special case of the GARCH(1,1) model by enforcing an assumption of conditional normality for the series being modeled. This assumption of conditional normality is problematic, however, as even conditional stock returns are generally found to be leptokurtic.[5] In addition, the model is specified only for the simple GARCH(1,1) form of conditional volatility. Gray does indicate (p. 32): "Although not investigated in this paper, the GRS framework can be extended to accommodate other conditional variance processes." He does not, however, provide specifics of how this would be accomplished.

Baldauf and Santoni (1991) investigate the hypothesis that the advent of program trading has caused a shift in the level of volatility in the underlying stock market. They propose a model in which the parameters of the ARCH conditional volatility equation are allowed to change at a discrete point:

$$h_t^2 = \omega + \sum_{i=1}^{q} \alpha_i u_i^2 + \beta_0 D_f + \sum_{i=1}^{q} \beta_i u_i^2 D_f \qquad (1.27)$$

In this specification, D_f is a $(0,1)$ dummy variable which is zero prior to the date corresponding to the advent of program trading and one otherwise. They find no evidence of a change in the parameters of the ARCH equation corresponding to the advent of program trading. They do note, however, (p. 198): "Identifying the advent of program trading is problematic."

These approaches to combining the evolutionary nature of GARCH models with models that allow for changing unconditional volatility provide potential solutions in specific situations. A more general framework for determining whether the two methods of modeling time varying volatility provide different descriptions of the same phenomenon or whether they model different concurrent phenomena is developed and tested in the following chapters.

The remainder of this study is organized as follows: Chapter 2 considers topics related to estimation of the highly nonlinear and frequently poorly behaved objective functions which are associated with models based on the GARCH framework.

---

[5] (See, for example, Campbell et al., 1997, pp. 488–489)

Chapter 3 takes the Nested GARCH model of Hentschel (1995) as a starting point, and determines the most appropriate parameterization of the model for subsequent enhancement and development. Chapter 4 continues the development of the particular parameterization of the Nested GARCH model. Alternative specifications for the conditional mean portion of the model are considered first, following which the Regime Dependent GARCH (RD-GARCH) model is introduced. Chapter 5 then subjects the RD-GARCH model to a number of tests to determine how well it performs at describing the volatility of an U.S. equity index over a particular sample period. Chapter 6 continues the evaluation of the RD-GARCH model, investigating the stability and robustness of the results obtained in Chapter 5. Chapter 7 concludes the evaluation of the RD-GARCH methodology with an investigation of how well it performs in predicting future (out-of-sample) variance. Finally, Chapter 8 provides a general summary and concluding remarks.

Chapter 2

Estimation of GARCH Models

In general the estimation of highly non-linear models, which includes several members of the GARCH family, poses several difficult problems. In particular, the models estimated for this study require the maximization of a recursive likelihood function of up to 23 parameters. Some of the parameter restrictions imposed on the nested GARCH conditional variance equation make the likelihood function poorly behaved—either discontinuous, or extremely noisy over the range of the remaining parameters. This chapter considers two topics related to the estimation of GARCH models: the optimization of nonlinear objective functions, and differentiation of the objective function with respect to its individual parameters.

2.1 Optimization

The purpose of numerical optimization is to determine the combination of parameters which provide the largest (maximization) or smallest (minimization) value for an objective function. The choice of whether maximization or minimization is used is largely arbitrary, as it is usually possible to change a maximization problem into a minimization problem simply by multiplying the objective function by -1.[1]

One of the most important distinctions in numerical optimization is the nature of the solution provided. Underlying many of the more common numerical optimization techniques is an implicit assumption that the function to be maximized is concave, or very nearly so, over some relatively predictable region. The optimization routine is initialized to some plausible starting point and moves in some (usually deterministic) manner to the point where the value of the objective function is largest. These techniques work well for functions which are strictly concave since any movement resulting in a larger value for the objective function is a move towards the solution of the optimization problem. Choice among alternative techniques is then motivated by how well a particular technique handles the shape of the specific objective function.

---

[1] In fact, the techniques chosen for this study solve a minimization problem rather than a maximization problem because of the availability of an "off the shelf" minimization routine which performs well. In the discussion, reference to maximization of the log-likelihood function should be considered to be equivalent to minimization of the negative of the function.

For example, some techniques find the solution more quickly for objective functions characterized by a steep peak in the vicinity of the solution, while other techniques perform well for objective functions which are characterized by a more gentle slope near the solution.

If the objective function is not strictly concave, however, a problem may arise. If an optimization technique which operates by starting at a particular point and always moving in the direction which results in a larger value for the objective function finds a local maximum, it may well stop at that point and report it as the (possibly incorrect) solution. Techniques working in this manner are usually referred to as local optimization techniques since the solution that they provide is frequently the local maximum closest to the starting point. This behaviour may not be problematic if the function to be optimized is reasonably well behaved, and if it is possible to determine an appropriate starting point. If the function is very noisy or if there is no good reason to specify a particular starting point, however, local optimization techniques become difficult to apply.[2]

The alternatives to local optimization techniques are referred to as global optimization techniques. These techniques are much less common than local optimization techniques at least partly because the problem of global optimization is in principle impossible to solve without an exhaustive search of the function's parameter space. Such a search would, unfortunately, not be possible to perform in a finite amount of time.[3] The problem of global optimization has led in part to the distinction between *deterministic* and *stochastic* optimization techniques because a deterministic solution to the problem of locating the global maximum of a general non-concave objective function does not appear to exist.

2.1.1 Deterministic vs. Stochastic Methods

In general, deterministic techniques (also commonly known as deterministic hill-climbing techniques[4]) take as given a starting point and choose, in some deterministic manner, the direction and step size from that point which makes the largest improvement in the function value. The point located in this manner becomes the new starting point and the process is repeated until no further significant improvement is possible. The difference between techniques is largely based on the method of choosing direction, step size, or both. The BHHH algorithm of Berndt et al. (1974) is the most common technique of this type in use in the field of econometrics, although there are several others. In particular, this study uses the DMNGB algorithm from the PORT subroutine library (Phyllis A. Fox, 1984) based on the TOMS 611 algorithm of Gay

---

[2] One alternative is to start the local optimization routine from several different randomly chosen points in the parameter space, and to choose the best solution (the solution which provided the largest likelihood value) as the global optimum. This technique does require that a large number of potential starting points be investigated for reasonable assurance of determining the global optimum. As the time required to find the local optimum from each starting point can be in the order of several hours, a search for the global optimum using this technique would require a prohibitive amount of computer time.

[3] In fact it is not possible to perform an exhaustive search in an infinite amount of time.

[4] This is a somewhat unfortunate misnomer in the case of minimization.

(1983). This algorithm was chosen over the BHHH algorithm because the PORT routine tended to produce results which were as good as or better than the results obtained from the BHHH algorithm,[5] while requiring significantly less computer time.

The operation of stochastic optimization techniques are generally less obvious than the operation of deterministic techniques. Two different stochastic techniques were investigated for this study: Simulated Annealing and a Genetic or Evolutionary algorithm.

The technique known as *simulated annealing* is conceptually similar to the group of deterministic hill-climbing techniques. A "current" point (analogous to the starting point in the deterministic technique) is randomly chosen by the algorithm. The algorithm then chooses a random step size and direction, computes the value of the objective function at the new point, and compares the new point to the current point. If the new point results in a larger value for the objective function, the new point becomes the current point and the process repeats. A new point which results in a smaller value for the objective function may be selected as the current point, however, with a probability selected by the implementor of the algorithm. This provides a mechanism for the algorithm to escape from a local maximum. The (random) step size is controlled by a parameter of the algorithm which is set to a large value at the beginning of the optimization process, and is reduced as the procedure progresses. This allows the algorithm to search widely over the parameter space in the beginning, but as the search progresses the parameter is refined and the range reduces. The rate at which the step-size parameter is reduced is controlled by another parameter of the algorithm. The process is stopped when the maximum value located for the objective function has not improved for a specified number of iterations. The particular choice of values for the parameters of the algorithm: the probability of accepting a smaller current point, the initial step size, the rate of reduction of the step size, and the stopping criteria is somewhat ad-hoc, and based upon practice and experience with the technique. An algorithm based on this technique (see Goffe et al., 1994) was used in early work for this study, but was abandoned due to the computer resources required.[6]

A second category of solutions to the global optimization problem are known as *genetic or evolutionary algorithms* and encompass a large variety of related techniques. The common feature of these methods is that a "population" of many potential solutions is maintained, and at each "generation" of the algorithm new potential solutions are created by altering one or more parameter values in each member of the existing population. The potential new member of the population replaces a member of the existing population if it provides a better solution. These techniques differentiate themselves from each other in the method used to select and alter (or perturb) parameter values. The suitability of a particular technique for a particular optimization problem is related to how well the selection and perturbation methods fit with the type of parameters and nature of parameter restrictions which comprise the model to be optimized.[7] The method chosen for this study is the Differential Evolution (DE) technique of Price and Storn (1997). This technique works well with real-valued parameters, and handles the bounds constraint restrictions which the parameters of these models require. In addition, the technique is simple to implement and performs well relative to other related techniques.[8] The DE algorithm works as follows:

---

[5] The likelihood values obtained from the PORT algorithm were as large as or larger than those obtained from the BHHH algorithm given identical starting points.

[6] An early test of a simplified version of the model required over 2.3 million function evaluations, and it took several days of computer time to arrive at a solution.

1. An initial population of possible solutions is generated. The size of this population is ten times the number of parameters to be estimated; if 17 parameters are to be estimated, a population of 170 possible solutions is initialized.[9]

2. For each member of the population, new values for one or more parameters are generated by taking a linear combination of the same parameter from two other members of the population. Both the particular parameters which are generated in this way and the population members from which the linear combination is constructed are randomly selected from a uniform distribution. The values obtained for the new parameters are subject to the same range checks outlined for Step 1.

3. If the value of the objective function resulting from the changes to the parameter values is larger than the value from the un-perturbed population member, the new member replaces the old member.

4. The algorithm returns to Step 2 and continues until the appropriate stopping criteria is met (usually a number of iterations, or generations in which there have been no changes to the best member of population).

As with simulated annealing, genetic algorithms tend to be ad-hoc in both the selection of control parameters and in the specification of stopping criteria. They are also similar to simulated annealing as they converge slowly. They do have significant advantages over simulated annealing, however:

- As the population evolves, it provides a map of the likelihood surface. While the applications of this map is not considered in this study, it has potential benefits for future development.

---

[7] The perturbation technique used with some genetic algorithms assumes that the model parameters are strings of $(0, 1)$ binary digits. With others, the model parameters are assumed to take on only integer values, or to be strings of characters, etc. While it is in principle possible to use a technique which assumes that the parameters are strings of binary digits to optimize a function of real-valued parameters, such an approach would likely perform poorly.

[8] The Differential Evolution technique finished 3rd at the First International Contest on Evolutionary Computation held at Nagoya, Japan, in May 1996. The two algorithms which outperformed DE did solve the test problem more quickly, but were not general purpose techniques.

[9] Each of these possible solutions is generated using series of uniformly distributed random numbers bounded by the range of possible values for the various parameters. For example, if this algorithm were used to optimize a likelihood function based on the GARCH(1,1) model defined by Equation 1.7, the population would be initialized with values for α and β uniformly distributed in the range (0,1). An additional restriction that $\alpha + \beta < 1$ would also be imposed to ensure stationarity.

- The generation of new potential population members (Step 2 above) may be executed in parallel allowing the problem to be solved more quickly on either a parallel computer or a cluster of workstations.
- The best population member from a generation prior to the final solution can be used as a plausible starting point for a much faster deterministic technique.[10]

The estimation of the models constructed in this study have been performed by a combination of stochastic and deterministic techniques. The stochastic technique is executed first, and used to locate a likely starting point for the deterministic technique. Once the deterministic technique completes, the stochastic technique is employed again to ensure that the solution is a global optimum. The complete optimization algorithm is as follows:

1. The DE algorithm is executed for 100 generations.
2. The parameter estimates from best potential model are used as a starting point for the PORT deterministic algorithm.
3. The solution from the PORT algorithm replaces the best potential solution from the DE algorithm, and the DE algorithm is restarted.
4. Execution of the DE algorithm continues for an additional 100 generations.
5. If there has been a change in the best potential solution, return to Step 2.
6. If there has been no improvement in 100 generations of the DE algorithm, stop and report solution.

This algorithm has several advantages: it does not require the specification of an initial starting point or guess; Step 4 will allow an escape from a local optimum if the deterministic portion of the algorithm should reach such a point; it will also converge much more quickly than a technique which is based solely on the DE algorithm since deterministic techniques generally converge much more quickly than stochastic techniques. It does suffer from the ad-hoc nature of the stochastic portion of the algorithm. The selection of 100 as the number of iterations of the DE algorithm for locating a starting point for the PORT algorithm was based on empirical tests, as was the selection of 100 as the number of iterations with no change in best potential solution for the stopping criterion. While these values work well in practice for the particular set of models used in this study, there is no method for formally determining the optimal values for these parameters. Similarly, there is no way to determine if the point identified as the global optimum by this technique is in fact the global optimum. As will be shown later,[11] however, this combined technique outperforms the alternative of using a plausible starting point for the deterministic algorithm.

---

[10] It is also possible to use an deterministic algorithm such as Nelder and Mead (1965) which is less sensitive to its starting point to determine a plausible starting point for the more sensitive but quicker hill-climbing technique. This is the approach used by RATS (Estima, 1995), but this procedure is only slightly better than simply using a conventional deterministic hill climbing technique alone as it still requires specification of *some* starting point.

[11] Table 3.1 in Section 3.1.3

2.2 Derivatives

A common requirement for deterministic hill-climbing optimization techniques is that a method be provided for determining the value of the function to be optimized, its first derivative, and possibly higher order derivatives, at arbitrary points in the function's parameter space. These values are then used by the algorithm to determine the direction and distance (step size) to the next test point. In addition, to be able to conduct statistical inference on the parameter estimates produced by any optimization procedure, it is necessary to compute some measure of the standard errors for the parameter estimates. In the case of maximum likelihood estimation, it is possible to obtain the asymptotic covariance matrix of the parameter estimates using the first derivatives of the likelihood function with the BHHH estimator of the information matrix as suggested by Berndt et al. (1974):

$$\mathfrak{I}_{OPG} = \sum_t \boldsymbol{g}_t \boldsymbol{g}_t'$$

where

$$\boldsymbol{g}_t = \frac{\partial \mathfrak{L}_t(r_t, \gamma)}{\partial \gamma}$$

where $\mathfrak{L}_t(r_t, \gamma)$ represents the contribution to the log-likelihood of the t^{th} observation. The asymptotic estimator of the parameter covariance matrix is then:

$$\boldsymbol{V}_{OPG} = \mathfrak{I}_{OPG}^{-1}$$

which is the inverse of the outer product of the gradient of the log-likelihood function, and is known as the OPG estimator. The OPG estimator is asymptotically valid when the likelihood function is properly specified.

It is common to assume normally distributed disturbances when constructing the likelihood equation for models from the GARCH family. This is problematic, however, as it implies that the distribution of stock returns is conditionally normal and that the distribution of the standardized residuals should be unconditionally normal. In practice, however, returns exhibit excess kurtosis, and these "fat tails" lead to misspecification of the likelihood equation under the assumption of normality (see Campbell et al., 1997, pp. 488–489). If the distributional assumption is relaxed merely to allow for finite 4^{th} moments, and the likelihood equation obtained under the assumption of normality is used, the parameter estimates obtained from the misspecified model will be consistent, but the covariance matrix must be adjusted. It is appropriate in this case to use the quasi maximum likelihood (QML) estimator of Bollerslev and Woolridge (1992) to compute the covariance matrix:

$$\boldsymbol{V}_{QML} = \boldsymbol{H}^{-1} \mathfrak{I}_{OPG} \boldsymbol{H}^{-1}$$

where

$$H = \frac{\partial^2 \mathcal{L}(\boldsymbol{r}, \boldsymbol{\gamma})}{\partial \boldsymbol{\gamma} \partial \boldsymbol{\gamma}'}$$

Is the Hessian of the likelihood function. Thus, for the models considered in this study, it is necessary to compute both the first and second derivatives of the log-likelihood function. There are, in general, three ways that this can be done: analytic derivatives, numeric derivatives, and automatic differentiation. These techniques will be discussed in the following three sections.

2.2.1 Analytic Derivatives

An intuitively attractive method for computing the partial derivatives of the log-likelihood function would be to derive the functional form of each derivative and to code it directly into the optimization program. Such a procedure generally results in derivatives which are both very accurate, and require a minimum of computer time to compute at arbitrary points in the parameter space. In many cases, however, this functional form is difficult or impossible to derive.

2.2.2 Numeric Derivatives

It is also possible to use only evaluations of the objective function to compute numerical approximations to the derivative. The standard method of computing such approximations is to start with the definition of the first derivative:

$$\frac{\partial f(x)}{\partial x} = f'(x) = \lim_{h \to 0} \frac{f(x+h) - f(x)}{h} \tag{2.1}$$

and to replace h in this formulation by some small value, δ. This approach yields non-central finite-difference numeric derivatives which are a popular choice because they seem intuitively attractive, and require only one additional evaluation of the objective function for each parameter of the objective function. There are two significant problems with this approach, however:

- Rounding errors introduced by the finite nature of the internal representation of real numbers in computers are a problem when performing a series of computations. This problem becomes much more pronounced when the numbers being operated on differ significantly in magnitude. As the value δ must be much smaller than the value of $f(x)$ for Equation 2.1 to provide even a plausible estimate of $f'(x)$, extreme care must taken in selecting a value for δ. If this is not done, the rounding errors will accumulate quickly and render the estimate of $f'(x)$ meaningless. It is possible to reduce this problem by using central finite difference numeric derivatives which do not require as small a value for δ:

$$f'(x) = \lim_{h \to 0} \frac{f(x+h) - f(x-h)}{2h} \tag{2.2}$$

This form provides better estimates of the derivative of the function, at the cost of additional function evaluations. Even better estimates can be obtained by computing additional values of the objective function in the vicinity of the point at which the derivative is required and using quadratic or cubic (or higher order) interpolation. These methods do not entirely remove the problem of rounding errors, however. In addition, it becomes very costly in computer time to compute derivatives using these methods as the number of function evaluations rises quickly.

- A second and more serious problem results from the nature of the approximation. Equation 2.1 is valid *only* in the $\lim_{h \to 0}$. For the case where $h = \delta$, the Taylor series expansion

  $$f(x+\delta) = f(x) + \delta f'(x) + \frac{\delta^2}{2!} f''(x) + \frac{\delta^3}{3!} f'''(x) + \ldots$$

  applies. Rearranging the above to

  $$f'(x) = \frac{f(x+\delta) - f(x)}{\delta} - \frac{\delta}{2!} f''(x) + \frac{\delta^2}{3!} f'''(x) + \ldots$$

  and using Equation 2.1 with $h = \delta$ to compute numeric first derivatives implicitly truncates the Taylor series at the second term. While this approximation performs reasonably well for functions which have well-behaved first derivatives,[12] attempting to apply Equation 2.1 twice to obtain the second derivative to compute the covariance matrix is pointless; the first application has made the implicit assumption that the second derivative is unimportant—it is meaningless to make a second application to obtain the Hessian.

2.2.3 Automatic Differentiation

A third method is to apply an automatic differentiation algorithm directly to the computer program which computes the function value. In essence the automatic differentiation program examines the program or subroutine which computes the function value and produces a second program which computes the derivative of that program or subroutine. It does this by computing the derivative of each line of the original program and applying the chain rule to these derivatives of individual program statements to provide a routine which computes the derivative of the program. Since the derivative program which is produced in this way computes the analytic derivative rather than a numeric approximation, the automatic differentiation program can be applied to the first derivative program to produce a third program which generates second derivatives.

---

[12] This unfortunately does not appear to be the case for models considered in this study.

This study uses the ADIFOR program of Bischof et al. (1995) to produce programs which generate both the first and second derivatives of the likelihood function. The results obtained from these programs for first derivatives were compared to the values obtained from analytic derivatives generated by the Maple V (Waterloo Maple Inc., 1998) symbolic mathematics system and to the results from both first and third order central finite-difference numeric differentiation, and were found to be nearly identical.

Chapter 3

Model Selection

3.1 Replication of Previous Results

Because it nests many of the more popular GARCH models into a single equation, the Nested GARCH model of Hentschel (1995) forms an excellent foundation for the extensions to the GARCH framework which comprise this study. As such, a natural starting point is with the replication of his results. This serves two useful purposes:

1. His work did not report several parameter values which were incidental to his work but which are of interest here.
2. He did not report the likelihood values obtained from the maximum likelihood estimation of his models. These values are necessary here to assess whether the model enhancements proposed in this study provide improvements over conventional GARCH models.

3.1.1 Data

The data used for the model selection and extension portions of this study comprise daily excess returns on U.S. equities covering the period January 2, 1926 through December 31, 1990. The daily stock returns for the period January 2, 1926 through July 2, 1962 are from Schwert (1989), and for the period July 3, 1962 through December 31, 1990 are from the Center for Research in Security Prices (CRSP) value weighted stock indices for NYSE and AMEX combined. The risk-free interest rate data are the one month average rate from the Fama/Bliss dataset in the CRSP bond data product.

Following Hentschel (1995), the logged monthly annual rate from the dataset was divided by 365 to give an estimate of the daily risk-free rate. That is:

$$r_t = \ln(1 + R_t) - \frac{\ln(1 + R_{fm}/100)}{365} \tag{3.1}$$

Where R_t is the daily raw return series from either the Schwert or the CRSP dataset as appropriate and R_{fm} is the monthly risk-free rate from the Fama/Bliss dataset for the month corresponding to the daily date t.

3.1.2 Model

Hentschel's work assumes a GARCH-M mean equation:

$$r_t = \gamma_0 + \gamma_1 h_t^2 + u_t \tag{3.2}$$

where u_t has zero mean and conditional variance h_t^2. As noted earlier, the standardized residual $v_t = (u_t/h_t) \sim iid(0,1)$. The conditional variance equation takes the form:

$$\frac{h_t^\lambda - 1}{\lambda} = \omega' + \alpha' h_{t-1}^\lambda f^\nu(v_{t-1}) + \beta \frac{h_{t-1}^\lambda - 1}{\lambda} \tag{3.3}$$

where

$$f(v_t) = |v_t - b| - c(v_t - b) \tag{3.4}$$

The form of Equation 3.3 is not convenient for estimation purposes. As outlined in Section 1.2.5, it may be rewritten as

$$h_t = \left(\omega + \alpha h_{t-1}^\lambda f^\nu(v_{t-1}) + \beta h_{t-1}^\lambda\right)^{\frac{1}{\lambda}} \tag{3.5}$$

where

$$f(v_t) = |v_t - b| - c(v_t - b)$$

for the case where $\lambda \neq 0$, and for the case where $\lambda = 0$ as

$$h_t = \exp(\omega + \alpha f^\nu(v_{t-1}) + \beta \log h_{t-1}) \tag{3.6}$$

which represents one form of the Exponential GARCH specification. A second issue associated with the numerical optimization of this function involves the $f(v_t)$ term. As this term involves an absolute value, it is not differentiable. The solution adopted mirrors that used by Hentschel, where the absolute value term is replaced:

$$f(v_t) = |v_t - b| - c(v_t - b) = \sqrt{a^2 + (v_t - b)^2} - c(v_t - b) \tag{3.7}$$

As a goes to zero the $\sqrt{a^2 + (v_t - b)^2}$ term converges to $|v_t - b|$. Following Hentschel's suggestion, the value $a = 0.001$ is used. For the case of normally distributed disturbances, these specifications result in a log-likelihood function of the form:[1]

$$\mathcal{L} = \sum_{t=1}^{T} \left[-\frac{1}{2} \log 2\pi - \log h_t - \frac{1}{2} v_t^2 \right] \tag{3.8}$$

Table 3.1: Comparison of estimation techniques

Model	with DE	without DE	Difference
model01	58628.527	58628.527	0.000
model02	58780.175	58780.175	0.000
model03	58781.598	58781.598	0.000
model04	58784.379	58784.379	0.000
model05	58608.295	58608.295	0.000
model06	58763.779	58763.779	0.000
model07	58775.111	58775.111	0.000
model08	58775.805	58775.805	0.000
model09	58643.384	58643.384	0.000
model10	58797.149	58797.149	0.000
model11	58759.586	58759.586	0.000
model12	58797.183	58797.183	0.000
model13	58651.975	58651.824	0.151
model14	58807.403	58781.545	25.858
model15	58793.483	58761.670	31.813
model16	58808.688	58808.663	0.024
model17	58655.439	58655.439	0.000
model18	58810.909	58810.909	0.000
model19	58794.833	58794.833	0.000
model20	58811.752	58811.752	0.000

The table presents the likelihood values obtained with two different numerical estimation techniques. The first column shows the likelihood value which results when the likelihood function is maximized by a combination of the (stochastic) DE algorithm and the deterministic hill-climbing technique. The second column shows the value which results when a plausible starting point was chosen for the deterministic hill-climbing technique. The starting point chosen was: $h_0 = 0.01$, $\lambda = \nu = 2.0$, $\alpha = 0.1$, and $\beta = 0.8$. All other parameters were set to zero. For the models where λ and ν are restricted to be some value other than 2.0, the appropriate restricted value was used. The model names "model01" through "model20" refer to the 20 models corresponding to the exhaustive list of model definitions possible with the parameter restrictions suggested by Hentschel (1995). A summary of the composition of the models is presented in Table A.1 in Appendix A. The data set comprises excess returns on an U.S. equity index covering the period January 2, 1926 to December 29, 1990, a total of 17,484 observations.

3.1.3 Results

A total of twenty sub-models of the nested GARCH framework were estimated. Table 3.1 reports the likelihood values obtained with two different numerical techniques for estimating these sub-models, and Table A.1 in Appendix A provides a summary of the parameter restrictions under which the models were constructed. In most cases, the two alternatives result in identical likelihood values at the optimum. For the models which impose the restriction $\lambda = \nu$ (model13 through model16), however, the combined technique yields a superior solution. In one case the likelihood ratio test for model selection conducted on the results obtained with the hill-climbing technique alone results in negative likelihood ratios. This fact, which clearly illustrates the difficulties and dangers inherent in manually selecting a starting point for deterministic optimization of nonlinear models,[2] led to the use of the combined technique in spite

---

[1] For a derivation of the log-likelihood function, see Appendix C.1

[2] See Section 2.1

Table 3.2: Estimates of model parameters

Model	γ_0	γ_1	λ	ν	b	c	ω	α	β	Likelihood
Models where $\lambda = 0, \nu = 1$ (Exponential GARCH)										
model01	$4.206^{[-4]}$ $(8.580^{[-5]})^*$	2.183 $(0.912)^*$	0	1	0	0	-0.131 $(0.016)^*$	0.101 $(0.011)^*$	0.988 $(0.002)^*$	58628.527
model02	$4.031^{[-4]}$ $(1.110^{[-4]})^*$	-0.582 (1.290)	0	1	0.424 $(0.043)^*$	0	-0.165 $(0.019)^*$	0.104 $(0.010)^*$	0.983 $(0.003)^*$	58780.175
model03	$3.292^{[-4]}$ $(8.116^{[-5]})^*$	0.524 (0.981)	0	1	0	0.362 $(0.045)^*$	-0.127 $(0.016)^*$	0.084 $(0.008)^*$	0.987 $(0.002)^*$	58781.598
model04	$3.760^{[-4]}$ $(8.222^{[-5]})^*$	-0.189 (1.027)	0	1	0.320 $(0.016)^*$	0.114 $(0.044)^*$	-0.156 $(0.018)^*$	0.096 $(0.009)^*$	0.984 $(0.003)^*$	58784.379
Models where $\lambda = 1, \nu = 1$ (Absolute Value GARCH)										
model05	$4.352^{[-4]}$ $(8.766^{[-5]})^*$	1.933 $(0.923)^*$	1	1	0	0	$1.032^{[-4]}$ $(2.058^{[-5]})^*$	0.106 $(0.012)^*$	0.911 $(0.010)^*$	58608.295
model06	$4.276^{[-4]}$ $(7.881^{[-5]})^*$	-0.966 (0.827)	1	1	0.430 $(0.031)^*$	0	$1.523^{[-4]}$ $(2.568^{[-5]})^*$	0.112 $(0.011)^*$	0.892 $(0.011)^*$	58763.779
model07	$3.673^{[-4]}$ $(7.344^{[-5]})^*$	-0.120 (0.785)	1	1	0	0.391 $(0.046)^*$	$1.263^{[-4]}$ $(2.421^{[-5]})^*$	0.092 $(0.010)^*$	0.918 $(0.009)^*$	58775.111
model08	$3.556^{[-4]}$ $(1.107^{[-4]})^*$	-0.213 (1.397)	1	1	0.030 (0.095)	0.371 $(0.097)^*$	$1.294^{[-4]}$ $(3.060^{[-5]})^*$	0.092 $(0.010)^*$	0.916 $(0.012)^*$	58775.805

Entries of the form $x^{[-y]}$ represent scientific notation; that is $2.5^{[-2]}$ represents the value 2.5×10^{-2}.

Table 3.2: continued.

Model	γ_0	γ_1	λ	ν	b	c	ω	α	β	Likelihood
Models where $\lambda = 2, \nu = 2$ (Conventional Value GARCH)										
model09	$4.523^{[-4]}$ $(7.890^{[-5]})^*$	1.574 $(0.733)^*$	2	2	0	0	$8.600^{[-7]}$ $(1.740^{[-7]})^*$	0.095 $(0.011)^*$	0.901 $(0.011)^*$	58643.384
model10	$4.300^{[-4]}$ $(8.046^{[-5]})^*$	−1.358 (0.903)	2	2	0.475 $(0.056)^*$	0	$1.081^{[-6]}$ $(2.289^{[-7]})^*$	0.087 $(0.010)^*$	0.887 $(0.013)^*$	58797.149
model11	$4.159^{[-4]}$ $(7.990^{[-5]})^*$	−0.539 (0.842)	2	2	0	0.256 $(0.031)^*$	$9.395^{[-7]}$ $(1.989^{[-7]})^*$	0.078 $(0.009)^*$	0.908 $(0.010)^*$	58759.586
model12	$4.305^{[-4]}$ $(8.067^{[-5]})^*$	−1.362 (0.904)	2	2	0.487 $(0.087)^*$	−0.009 (0.047)	$1.084^{[-6]}$ $(2.291^{[-7]})^*$	0.088 $(0.011)^*$	0.886 $(0.014)^*$	58797.183
Models where $\lambda = \nu$										
model13	$4.547^{[-4]}$ $(8.065^{[-5]})^*$	1.681 $(0.789)^*$	1.646 $(0.417)^*$	1.646	0	0	$4.669^{[-6]}$ $(1.005^{[-5]})$	0.102 $(0.028)^*$	0.905 $(0.021)^*$	58651.975
model14	$4.163^{[-4]}$ $(8.337^{[-5]})^*$	−1.067 (0.940)	1.614 $(0.358)^*$	1.614	0.455 $(0.051)^*$	0	$7.329^{[-6]}$ $(1.307^{[-5]})$	0.100 $(0.027)^*$	0.889 $(0.025)^*$	58807.403
model15	$3.852^{[-4]}$ $(8.748^{[-5]})^*$	−0.334 (0.978)	1.362 $(0.358)^*$	1.362	0	0.337 $(0.044)^*$	$2.081^{[-5]}$ $(3.680^{[-5]})$	0.089 $(0.028)^*$	0.916 $(0.022)^*$	58793.483
model16	$4.116^{[-4]}$ $(8.385^{[-5]})^*$	−1.000 (0.945)	1.569 $(0.365)^*$	1.569	0.377 $(0.070)^*$	0.067 (0.052)	$8.957^{[-6]}$ $(1.627^{[-5]})$	0.096 $(0.026)^*$	0.894 $(0.025)^*$	58808.688

Entries of the form $x^{[-y]}$ represent scientific notation; that is $2.5^{[-2]}$ represents the value 2.5×10^{-2}.

Table 3.2: continued.

Model	γ_0	γ_1	λ	ν	b	c	ω	α	β	Likelihood
Models where λ, ν are free										
model17	$4.452^{[-4]}$	1.823	0.915	1.518	0	0	$1.668^{[-4]}$	0.060	0.939	58655.439
	$(8.180^{[-5]})^*$	$(0.841)^*$	$(0.322)^*$	$(0.131)^*$			$(2.583^{[-4]})$	$(0.021)^*$	$(0.018)^*$	
model18	$3.863^{[-4]}$	−0.654	0.986	1.509	0.452	0	$1.558^{[-4]}$	0.064	0.924	58810.909
	$(8.377^{[-5]})^*$	(0.977)	$(0.292)^*$	$(0.118)^*$	$(0.054)^*$		$(2.189^{[-4]})$	$(0.019)^*$	$(0.020)^*$	
model19	$3.656^{[-4]}$	−0.047	0.949	1.303	0	0.336	$1.538^{[-4]}$	0.063	0.936	58794.833
	$(8.392^{[-5]})^*$	(0.961)	$(0.292)^*$	$(0.108)^*$		$(0.048)^*$	$(2.139^{[-4]})$	$(0.021)^*$	$(0.017)^*$	
model20	$3.848^{[-4]}$	−0.631	0.986	1.482	0.392	0.053	$1.533^{[-4]}$	0.062	0.925	58811.752
	$(8.577^{[-5]})^*$	(0.998)	$(0.293)^*$	$(0.114)^*$	$(0.085)^*$	(0.056)	$(2.159^{[-4]})$	$(0.019)^*$	$(0.020)^*$	

Entries of the form $x^{[-y]}$ represent scientific notation; that is $2.5^{[-2]}$ represents the value 2.5×10^{-2}.

The table presents the parameter estimates for nested GARCH models comprising the conditional mean specified by Equation 3.2 in the text, and parameter restrictions imposed on the conditional volatility specified by Equations 3.3 and 3.4 in the text. Numbers in parentheses under the parameter estimates represent asymptotic standard errors. Parameters which are significantly different from zero at the 5% level are indicated with an asterisk. A summary of the composition of the models is presented in Table A.1 in Appendix A. The data set comprises excess returns on an U.S. equity index covering the period January 2, 1926 to December 29, 1990, a total of 17,484 observations.

of the much higher computational time requirement.[3]

Table 3.2 reports the parameter estimates for the twenty sub-models of the Nested GARCH specification. These results are quantitatively similar and qualitatively identical to those reported by Hentschel with the exception of the model corresponding to Exponential GARCH.[4] The table is broken into five sections, with each section of the table corresponding to a general class of GARCH models. The first section of the table ($\lambda = 0, \nu = 1$) corresponds to the family of Exponential GARCH models. The second section ($\lambda = 1, \nu = 1$) represents Absolute Value GARCH models, which provide evolution of volatility through the square root of the conditional variance. The third section ($\lambda = 2, \nu = 2$) represents the "conventional" GARCH models which provide evolution of volatility through the conditional variance. The fourth and fifth sections ($\lambda = \nu$) and (λ, ν free) represent progressively less restricted models. Both of these final two sections nest Absolute value GARCH and conventional GARCH. The fifth section also nests Exponential GARCH while the fourth section does not. Within each section of the table, the first line shows the most restricted version of the model; the symmetric version which allows neither a shift nor a rotation of the NIC. The second and third lines show the parameter estimates for the models where the NIC is allowed to shift and and to rotate respectively, and the fourth line shows the estimates for the model where the NIC is allowed to both shift and rotate simultaneously. Asymptotic standard errors are shown in parentheses below the parameter estimates, and are calculated using the quasi-maximum likelihood estimator of the parameter covariance matrix as outlined in Section 2.2.

3.2 Model Selection

Table 3.3 shows the results for likelihood ratio tests of the form of the asymmetry in volatility. These results are similar quantitatively to those presented by Hentschel, and are identical qualitatively. The table is broken into five sections in the same

---

[3] Estimation of a conventional GARCH(1,1) model (model09) using the deterministic hill-climbing technique alone required approximately 70 seconds of computer time on a PentiumII 266 processor. Estimating the same model with the combined technique required approximately 2,000 seconds on the same computer system. Using a network of six Pentium based computers operating in parallel, the model was estimated using the combined technique in 800 seconds.

[4] The model reported here as similar to the EGARCH model (model03) is constructed with the parameter restrictions: $\lambda = 0$, $\nu = 1$, and $b = 0$ which results in the conditional volatility equation:

$$\ln h_t^2 = \omega + \alpha[|v_t| - cv_t] + \beta \ln h_{t-1}^2 \tag{3.9}$$

Nelson (1991), however, suggests the EGARCH conditional volatility equation:

$$\ln h_t^2 = \omega + \alpha[|v_t| - E|v_t| + cv_t] + \beta \ln h_{t-1}^2 \tag{3.10}$$

Computing an estimate of the term $E|v_t|$ requires a distributional assumption about the disturbances. As later sections of this study consider alternative distributions, the adjustment has not been made to the EGARCH specification reported here in order to maintain generality.

manner as Table 3.2. Within each section of the table, the first line shows the results for likelihood ratio tests of the null hypothesis of no asymmetry ($b = c = 0$) vs. the alternative of asymmetry in some form. The second line shows the results for tests of the null hypothesis of a NIC which is shifted ($c = 0, b$ free) vs. the alternative of a curve which is both shifted and rotated (b, c free), and the third line shows the results for tests of the null hypothesis of a NIC which is rotated ($b = 0, c$ free) vs. the alternative of a curve which is both shifted and rotated. The likelihood ratio statistics are all distributed as χ^2 with either one or two degrees of freedom. The corresponding p-values are included in parentheses.

In all cases, the symmetric model can be rejected in favour of an asymmetric model. In addition, in many cases one or more of the asymmetric models can be rejected in favour of asymmetry of a different form. In the final three sections of the table, the null hypothesis of a rotated NIC can be rejected in favour of the alternative of a NIC which is both shifted and rotated, but the null hypothesis of a shifted NIC cannot be rejected in favour of the alternative of a NIC which is both shifted and rotated; the shift of the NIC alone appears to be enough to explain the asymmetry. In the first two sections of the table, the null hypothesis of a rotated NIC cannot be rejected in favour of the alternative of a NIC which is both shifted and rotated, but the null hypothesis of a shift in the NIC can be rejected in favour of the alternative of a NIC which is both shifted and rotated; rotation of the NIC by itself appears to be sufficient to explain the asymmetry in these cases. Either a shift *or* a rotation of the NIC is sufficient to explain the asymmetry in the conditional volatility. A shift and a rotation are not required simultaneously for any of the models under consideration.

The left panel of Table 3.4 presents the model selection tests for models which impose the restriction $b = 0$, which specifies a NIC which is rotated but not shifted. This corresponds to the model suggested by the first and second sections of Table 3.3. The right panel presents the same tests for models which impose the restriction $c = 0$, which specifies a NIC which is shifted but not rotated. This corresponds to the model suggested by the third through fifth sections of Table 3.3.

The results for the rotated NIC indicate that both of the models corresponding to the sections of Table 3.3 which found the rotated NIC to be appropriate can be rejected in favour of one of the less restricted models, and by implication that a shift rather than a rotation of the NIC is the appropriate method of modeling the asymmetry in conditional volatility. The results for the shifted NIC indicate that all of the more restricted models can be rejected at conventional levels in favour of the model which does not impose restrictions on λ or ν. Together these results indicate that the appropriate model from this set of nested models is the model which imposes the restriction $c = 0$ and where λ, ν, and b are freely estimated. This corresponds to model18 in Table 3.2.[5] This model will be used as a starting point for development in subsequent chapters, and the conditional variance specification:

---

[5] Hentschel chose the even less restricted model where c is also freely estimated, model20 from Table 3.2. The choice of a different model is not evidence of different results, but rather represents a different interpretation of the same results.

Table 3.3: LR tests for asymmetry

		H_A		
Model form	H_O	$c = 0, b$ free	$b = 0, c$ free	b and c free
$\lambda = 0, \nu = 1$	$b = c = 0$	303.296 (< 0.001)	306.141 (< 0.001)	311.704 (< 0.001)
	$c = 0, b$ free			8.408 (0.004)
	$b = 0, c$ free			5.563 (0.018)
$\lambda = 1, \nu = 1$	$b = c = 0$	310.969 (< 0.001)	333.631 (< 0.001)	335.020 (< 0.001)
	$c = 0, b$ free			24.051 (< 0.001)
	$b = 0, c$ free			1.389 (0.239)
$\lambda = 2, \nu = 2$	$b = c = 0$	307.530 (< 0.001)	232.404 (< 0.001)	307.597 (< 0.001)
	$c = 0, b$ free			0.067 (0.795)
	$b = 0, c$ free			75.194 (< 0.001)
$\lambda = \nu$	$b = c = 0$	310.856 (< 0.001)	283.017 (< 0.001)	313.426 (< 0.001)
	$c = 0, b$ free			2.570 (0.109)
	$b = 0, c$ free			30.409 (< 0.001)
λ, ν free	$b = c = 0$	310.939 (< 0.001)	278.788 (< 0.001)	312.627 (< 0.001)
	$c = 0, b$ free			1.688 (0.194)
	$b = 0, c$ free			33.839 (< 0.001)

The table presents the results for likelihood ratio tests for the form of asymmetry in the conditional volatility. Within each section, the first line shows the results for likelihood ratio tests of the null hypothesis of no asymmetry ($b = c = 0$) vs. the alternative of asymmetry in some form. The second line shows the results for tests of the null hypothesis of a NIC which is shifted ($c = 0, b$ free) vs. the alternative of a curve which is both shifted and rotated (b, c free), and the third line shows the results for tests of the null hypothesis of a NIC which is rotated ($b = 0, c$ free) vs. the alternative of a curve which is both shifted and rotated. The likelihood ratio statistics are all distributed as χ^2 with either 1 or 2 degrees of freedom. The corresponding p-values are included in parentheses. The data set comprises excess returns on an U.S. equity index covering the period January 2, 1926 to December 29, 1990, a total of 17,484 observations.

Table 3.4: LR tests for model selection

	H_A			H_A	
H_O	$\lambda = \nu$	λ, ν free	H_O	$\lambda = \nu$	λ, ν free
$\lambda = 0, \nu = 1$		26.471	$\lambda = 0, \nu = 1$		61.467
		(< 0.001)			(< 0.001)
$\lambda = 1, \nu = 1$	36.745	39.445	$\lambda = 1, \nu = 1$	87.247	94.258
	(< 0.001)	(< 0.001)		(< 0.001)	(< 0.001)
$\lambda = 2, \nu = 2$	67.795	70.495	$\lambda = 2, \nu = 2$	20.508	27.519
	(< 0.001)	(< 0.001)		(< 0.001)	(< 0.001)
$\lambda = \nu$		2.700	$\lambda = \nu$		7.012
		(0.100)			(0.008)
$b = 0$ (rotated NIC)			$c = 0$ (shifted NIC)		

The table presents the results for likelihood ratio tests for the functional form of the GARCH model. The parameter restrictions $\lambda = 0, \nu = 1$ correspond to models from the EGARCH family. The parameter restrictions $\lambda = 1, \nu = 1$ correspond to models from the AGARCH family. The parameter restrictions $\lambda = 2, \nu = 2$ correspond to models from the conventional GARCH specification. The parameter restrictions $\lambda = \nu$ correspond to a family of models which nest the AGARCH and conventional GARCH specifications. The left panel presents the results for models which impose the restriction $b = 0$, which corresponds to models with a rotated but not shifted NIC. These tests are appropriate for the models suggested by sections 1 and 2 of Table 3.3. The right panel presents the results for models which impose the restriction $c = 0$ which corresponds to models with a shifted but not rotated NIC. These tests are appropriate for the models suggested by sections 3, 4, and 5 of Table 3.3. The likelihood ratio statistics are all distributed χ^2 with either 1 or 2 degrees of freedom. The corresponding p-values are included in parentheses. The data set comprises excess returns on an U.S. equity index covering the period January 2, 1926 to December 29, 1990, a total of 17,484 observations.

$$h_t = \left(\omega + \alpha h_{t-1}^{\lambda} |v_{t-1} - b|^{\nu} + \beta h_{t-1}^{\lambda}\right)^{\frac{1}{\lambda}} \tag{3.11}$$

will be referred to as the Extended GARCH model.

3.3 Residual Tests

As a final comparison of the various models under consideration, Table 3.5 provides additional information regarding the residuals from the models estimated. The table is broken into six sections, with the first section of the table reporting the statistical moments and the Jarque and Bera (1987) normality test for the residuals from the mean equation for model01,[6] and the following five sections of the table reporting the moments and Jarque-Bera statistic for the standardized residuals, first order autocorrelation of the square root of the conditional variance series, h_t (which represents the persistence of the conditional volatility), and the sign bias tests suggested by Engle and Ng (1993) for all twenty of the models. The composition of these latter sections of the table parallels Table 3.2.

---

[6] The residuals from the mean equations for the other 19 models are not reported as their moments do not differ qualitatively from those reported for model01.

Table 3.5: Descriptive statistics for standardized residuals

Model	Mean	Variance	Skewness	Kurtosis	Jarque-Bera	Persistence	Sign Bias	Negative	Positive
u_t	-4.294×10^{-4}	1.265×10^{-4}	−0.557	19.943	290630.693				
model01	-3.897×10^{-2}	0.999	−0.601	4.327	14690.760	0.96968	6.504	−7.424	−3.887
model02	-4.127×10^{-3}	1.000	−0.576	4.907	18505.923	0.96509	2.342	−4.420	−1.236
model03	-6.133×10^{-3}	1.000	−0.587	4.975	19032.058	0.96962	3.527	−4.445	−1.140
model04	-4.303×10^{-3}	1.000	−0.579	4.965	18932.327	0.96643	2.418	−4.263	−1.086
model05	-3.837×10^{-2}	0.999	−0.615	4.484	15754.171	0.97065	6.286	−7.665	−3.985
model06	-3.133×10^{-3}	1.000	−0.586	4.994	19168.968	0.96514	2.079	−4.662	−1.381
model07	-4.234×10^{-3}	1.000	−0.595	5.139	20271.443	0.96993	3.127	−4.296	−1.247
model08	-1.724×10^{-3}	1.000	−0.594	5.128	20185.422	0.96956	2.999	−4.269	−1.239
model09	-3.708×10^{-2}	0.999	−0.568	4.122	13317.372	0.96584	6.353	−6.106	−4.127
model10	-8.443×10^{-5}	1.000	−0.563	4.948	18755.522	0.96351	2.795	−3.107	−1.531
model11	-8.003×10^{-3}	1.000	−0.547	4.529	15812.757	0.96443	4.309	−3.616	−1.884
model12	-8.752×10^{-5}	1.000	−0.563	4.951	18784.500	0.96345	2.773	−3.110	−1.541
model13	-3.841×10^{-2}	0.999	−0.580	4.186	13747.565	0.96800	6.417	−6.498	−4.123
model14	-8.879×10^{-4}	1.000	−0.570	4.971	18947.490	0.96449	2.523	−3.433	−1.532
model15	-5.071×10^{-3}	1.000	−0.572	4.835	17981.224	0.96881	3.638	−3.898	−1.476
model16	-1.002×10^{-3}	1.000	−0.570	4.960	18870.434	0.96514	2.572	−3.431	−1.456
model17	-3.862×10^{-2}	0.999	−0.576	4.145	13483.209	0.96672	6.502	−6.678	−3.981
model18	-1.174×10^{-3}	1.000	−0.567	4.913	18522.509	0.96361	2.567	−3.555	−1.403
model19	-5.418×10^{-3}	1.000	−0.571	4.808	17787.954	0.96846	3.642	−3.991	−1.392
model20	-1.235×10^{-3}	1.000	−0.567	4.910	18502.969	0.96424	2.606	−3.547	−1.347

The table presents additional descriptive statistics for the residuals from the mean equation for the sub-models presented in Table 3.2. The first section presents the statistical moments for the residuals from the mean equation for model01, and the following five sections present the statistical moments and descriptive statistics for the standardized residuals from the each of the models estimated in this chapter. The statistical moments for the residuals from the mean equation for model02 through model20 do not differ significantly from those presented for model01 and are not presented here. The first through fourth columns present the first four statistical moments of either the residuals (first section) or the standardized residuals (all other sections) from the mean equation. The fifth column presents the Jarque and Bera (1987) test for normality. The sixth column presents the first order autocorrelation of the square root of the conditional variance of the standardized residuals, which is a measure of the persistence of the conditional volatility. The seventh through ninth columns present the results of the sign-bias tests suggested by Engle and Ng (1993), explained in the text. A summary of the composition of the models is presented in Table A.1 in Appendix A. The data set comprises excess returns on an U.S. equity index covering the period January 2, 1926 to December 29, 1990, a total of 17,484 observations.

The likelihood equation used for this section is derived under the assumption of normally distributed disturbances, implying that the residuals, u_t, should be distributed $N(0, h_t^2)$, and the standardized residuals, v_t, should be distributed $iiN(0, 1)$. For large samples[7] from the standard normal distribution, the standard deviation of the first four moments are given by:

$$\begin{aligned} SD_{\text{mean}} &= \sqrt{1/N} &&= 7.56 \times 10^{-3} \\ SD_{\text{var}} &= \sqrt{2/N} &&= 1.07 \times 10^{-2} \\ SD_{\text{skew}} &= \sqrt{6/N} &&= 1.85 \times 10^{-2} \\ SD_{\text{kurt}} &= \sqrt{24/N} &&= 3.70 \times 10^{-2} \end{aligned}$$

The sample means reported for the standardized residuals (those in the final five sections of the table) are within two standard deviations of the expected value of zero in fifteen of the twenty cases. In each of the five cases where the sample mean is more than two standard deviations away from zero, the sign of the mean is negative. Each of the models with a statistically significant negative sample mean for the standardized residuals corresponds to a model which imposes the restriction that the response of the conditional volatility to unexpected shocks is symmetric. It would appear that in these cases the model is unable to account for an asymmetry in volatility which exists in the return series. The sign is also negative in the fifteen cases where the sample mean is less than two standard deviations away from zero. Although these values are not statistically significantly different from zero, nor are the data sets independent, the fact that all fifteen sample means have a negative sign is worth noting. This would seem to indicate that the method adopted by these models for dealing with asymmetry reduces the residual asymmetry demonstrated in the sample mean to insignificant levels, but does not remove it entirely.

The sample variance reported for the standardized residuals is within two standard deviations of the expected value of one in all twenty cases. This indicates that the models all perform well at accounting for the volatility in the observed series.

Both the sample skewness and kurtosis for all twenty models are more than two standard deviations from the value of zero which would be expected if the standardized residuals were normally distributed. In all cases, the residuals exhibit significant negative skewness and significant excess kurtosis. This reinforces the observation that the distribution of stock returns is negatively skewed and fat-tailed, and indicates that the GARCH methodology, in its current form, is unable to account for this departure from normality. It is useful to compare the results for excess kurtosis reported in the first section of Table 3.5 for the non-standardized residuals to those presented in the following five sections for the standardized residuals. The results indicate that allowing for GARCH effects in the residuals has reduced the level of excess kurtosis, but that a significant level remains. It is also interesting to note that the level of skewness is slightly larger for each of the symmetric models (the first model in each section of the table) than it is for the asymmetric models. While skewness is statistically significant in all cases, the fact that it is slightly larger in models which enforce symmetric response of conditional volatility to unexpected shocks reinforces the observation made in respect to the sample mean.

---

[7] For the data set used in this section, $N = 17,484$

Since the evidence from the first four sample moments indicates that the normality assumption upon which the likelihood function is based is problematic, it is appropriate to perform a test for normality of the residuals and the standardized residuals. The sixth column of Table 3.5 presents the Jarque-Bera test statistic for normality of observations:

$$JB = N\left(\frac{\hat{u}_3^2}{6} + \frac{\hat{u}_4^2}{24}\right) \tag{3.12}$$

where $\hat{u}_3$ and $\hat{u}_4$ are the coefficients of skewness and excess kurtosis respectively, and N is the number of observations in the sample. Under the null hypothesis of normality, the test statistic is distributed $\chi^2(2)$ which provides a 5% critical value of 5.99. The null hypothesis of normality is rejected in all twenty-one cases. There is a large decrease in the value of the statistic from a value of $290,631$ for the non-standardized residual to a value of $14,691$ for the standardized residual from model01, which indicates that allowing for GARCH effects in the residuals has decreased the departure from normality observed. The null hypothesis of normality is still strongly rejected. This would tend to indicate that either the model remains improperly specified or that some other distributional assumption for the disturbances would be appropriate. These alternatives will be investigated in later sections of this study.

The seventh column of the table shows the first order autocorrelation of the square root of the conditional variance series, h_t, and indicates that the model form has little impact on the degree of persistence in the conditional volatility.

The final three columns report the result of the Engle and Ng (1993) sign bias tests. These tests are intended to detect any predictable impact of lagged shocks on the squared standardized residual series which has not already been impounded into the model. The existence of such a predictable impact would indicate that the model is incorrectly specified.

The Sign Bias test statistic (column eight) is the t-ratio from the coefficient η_1 in the regression:

$$v_t^2 = \eta_0 + \eta_1 S_{t-1}^- + \epsilon_t \tag{3.13}$$

where S_t^- is a $(0,1)$ dummy which takes on the value one if the residual u_t is negative and zero otherwise. Thus, this test indicates whether the sign of the previous residual contributes additional information not captured by the model under consideration. The negative and positive sign bias tests (columns nine and ten) are the t-ratios of the coefficient η_1 in the regressions:

$$v_t^2 = \eta_0 + \eta_1 S_{t-1}^- u_{t-1} + \epsilon_t \tag{3.14}$$

and

$$v_t^2 = \eta_0 + \eta_1 S_{t-1}^+ u_{t-1} + \epsilon_t \tag{3.15}$$

respectively, where S_t^- is as defined above and S_t^+ is a $(0, 1)$ dummy which is defined as $\left(1 - S_t^-\right)$. The tests indicate whether previous large and small shocks of the same sign have a different impact on the conditional variance which has not been captured by the model under consideration.

The tests indicate that, for all of the models under consideration, the sign of the previous shock retains some predictive power. The first model in each section of the table (the symmetric version of that particular class of models) has a much larger value for the test statistic, but the statistic remains significant at conventional levels for for all models. The results are similar for the negative sign bias test, indicating that the magnitude of previous negative shocks retains some predictive power for all models under consideration. Again, this test statistic is uniformly larger for the symmetric models than for the asymmetric models, but remains significant at conventional levels for all models. The positive sign bias test, however, provides evidence that the magnitude of previous non-negative shocks contains information which has not been predicted by the symmetric models, but this effect disappears when the model allows for the asymmetric response of conditional volatility to shocks.

Chapter 4

Model Extensions

Previous chapters have traced the development of the nested GARCH framework for modeling various forms of conditional heteroskedasticity in financial data, and have outlined some of the problems with estimation of models from this class. This chapter takes the Extended GARCH model selected at the end of Chapter 3:

$$h_t = \left(\omega + \alpha h_{t-1}^{\lambda}|v_{t-1} - b|^{\nu} + \beta h_{t-1}^{\lambda}\right)^{\frac{1}{\lambda}} \tag{4.1}$$

as a starting point and continues its development.

As outlined in Chapter 1, conditionally heteroskedastic models from the GARCH family comprise two parts: An equation which defines the evolution of a dependent variable in terms of one or more independent variables—which is the conditional mean of the series—and a separate equation which defines the evolution of the volatility of the residuals from the mean equation—which is the conditional volatility. This chapter considers several issues related to the extension of both the conditional mean specification and the conditional variance specification. Extensions to and alternative forms for the conditional mean specification are considered first, followed by extensions to the conditional volatility specification. In addition, alternative distributional assumptions for the residuals from the conditional mean equation are discussed. For simplicity, the terms mean equation and volatility equation refer to the conditional mean and conditional volatility equations from the GARCH models respectively.

4.1 Mean Equation Development

One alternative for the mean equation is provided in the GARCH in Mean or GARCH-M model suggested by Engle et al. (1987). In this model, the conditional variance enters directly into the mean equation as an independent variable, and is considered to be a measure of the risk premium in the market. The GARCH-M mean equation is:

$$r_t = \gamma_0 + \gamma_1 h_t^2 + u_t \tag{4.2}$$

where u_t has zero mean and conditional variance h_t^2. The parameter γ_1 is considered to be a measure of the risk premium in the sense that a positive value indicates a positive relationship between returns and conditional volatility; in essence, expected returns are higher in periods of high volatility than in periods of low volatility. Models of this form have been widely used (see, for example, Engle et al., 1987; French et al., 1987; Nelson, 1991; Glosten et al., 1993), although as Whitelaw (1994) notes (p.516):

> ... empirical investigations into the contemporaneous correlation between the first two moments of stock market returns yield decidedly mixed results ...

The GARCH-M specification will be used here as a starting point for development of the mean equation both because its use in Chapter 3 provides a useful standard against which to compare alternative specifications and because of its popularity in the finance literature.

4.1.1 Lagged Dependent Variables

While the choice of the GARCH-M specification for the conditional mean equation is popular, care must be taken in its application. Consider for example a simple zero mean AR(1) process:

$$y_t = \psi y_{t-1} + u_t \tag{4.3}$$

where $u_t \sim iid(0, \sigma_u^2)$ and $|\psi| < 1$ for stationarity. As the u_t's are independent, $\{y_t\}$ is a strictly linear process. If, however, an attempt were made to impose a misspecified model which did not include the y_{t-1} term, the misspecification would result in apparent non-linear dependence in the conditional variance of the series. To illustrate this, consider the conditional variance of y_t when treating its conditional mean as zero:

$$\begin{aligned} E\left(y_t^2 | I_{t-1}\right) &= E\left(\left[\psi^2 y_{t-1}^2 + 2\psi y_{t-1} u_t + u_t^2\right] | I_{t-1}\right) \\ &= \psi^2 y_{t-1}^2 + 2\psi y_{t-1} E(u_t) + E(u_t^2) \\ &= \psi^2 y_{t-1}^2 + \sigma_u^2 \end{aligned} \tag{4.4}$$

whereas its true conditional variance is:

$$E\left(y_t^2 | I_{t-1}\right) - E\left(y_t | I_{t-1}\right)^2 = \psi^2 y_{t-1}^2 + \sigma_u^2 - (\psi y_{t-1})^2 = \sigma_u^2 \tag{4.5}$$

Omitting y_{t-1} from the conditional mean equation introduces misleading results.

4.1.2 Weekend Effect

Another potential source of model misspecification is the so-called "weekend effect" first identified by French (1980) in which daily returns exhibit a smaller than average value on Mondays. This effect can be thought of as a simple but predictable shifting of the mean, which again causes apparent nonlinear dependencies in the conditional volatility if not accounted for in the model.

4.1.3 Augmented Conditional Mean Equation

By adding additional terms to Equation 4.2, it is possible to construct a mean equation that corrects for these possible sources of model misspecification. The general form of the augmented conditional mean equation is:

$$r_t = \gamma_0 + \gamma_1 h_t^2 + \sum_{i=1}^{N} \gamma_{i+1} r_{t-i} + \gamma_{N+2} D_M + u_t \tag{4.6}$$

where, as before, u_t has zero mean and conditional variance h_t^2, r_{t-i} represents the return lagged i periods, and D_M is a (0,1) dummy which is one if the observation r_t occurred on a Monday, and zero otherwise. As before, the standardized residual, $v_t \sim iid(0,1)$ is defined as:

$$v_t = \frac{u_t}{h_t} \tag{4.7}$$

For estimation, it is necessary to select a value for N, the number of lagged returns to include in the conditional mean specification. The approach employed here regresses the excess return on the daily stock index series described in Section 3.1.1 on the first three lags of the excess return and a constant:

$$r_t = \underset{(2.88)}{0.000248} + \underset{(4.88)}{0.0765}\, r_{t-1} - \underset{(2.09)}{0.0427}\, r_{t-2} + \underset{(0.27)}{0.00435}\, r_{t-3}$$

The t-statistics, shown in parentheses, were corrected with the Newey and West (1987) variance estimator for models with autocorrelated disturbances. As only the first and second lagged values of the return are significantly different from zero, the third lagged return was excluded from the final version of the augmented conditional mean equation:[1]

$$r_t = \gamma_0 + \gamma_1 h_t^2 + \gamma_2 r_{t-1} + \gamma_3 r_{t-2} + \gamma_4 D_M + u_t \tag{4.8}$$

---

[1] A likelihood ratio test was also performed comparing the unrestricted model which included all three lagged returns with the restricted models containing one and two lagged returns. The results of this test confirmed the decision to exclude the return lagged three periods from the conditional mean equation.

Equation 4.8 is an improvement over the simple GARCH-M specification of Equation 4.2 in that it allows for tests of model specification for the conditional mean equation. Likelihood ratio tests may be performed on the model to determine which of the explanatory variables are necessary and appropriate.

4.1.4 Likelihood Equation

To perform an ML estimation of this model, the log-likelihood equation must be obtained. Continuing with the development of the model initiated in Chapter 3, normally distributed errors are assumed.[2] Substituting Equation 4.7 into Equation 4.8, yields:

$$r_t = \gamma_0 + \gamma_1 h_{t-1}^2 + \gamma_2 r_{t-1} + \gamma_3 r_{t-2} + \gamma_4 D_M + v_t h_{t-1} \tag{4.9}$$

and upon rearranging

$$v_t = \frac{r_t - \gamma_0 - \gamma_1 h_{t-1}^2 \gamma_2 r_{t-1} + \gamma_3 r_{t-2} + \gamma_4 D_M}{h_{t-1}} = F_t(r_t, \gamma)$$

Following the procedure outlined in Appendix C.1 the log-likelihood is:

$$\mathcal{L} = \sum_{t=1}^{T} \left[-\frac{1}{2}\log 2\pi - \log h_t - \frac{1}{2} F_t(r_t, \gamma)^2 \right] \tag{4.10}$$

4.1.5 Mean Equation Selection

The third column of Table 3.2 presented the coefficients of the GARCH-M parameter obtained from the combination of the GARCH-M mean equation specification with the full range of Nested GARCH models. It is interesting to note that this coefficient is rarely significant, changing both in magnitude and sign frequently. If this parameter is viewed as a risk premium as Engle et al. (1987) suggest, its meaning is difficult to interpret in this setting.

Further inspection reveals that this parameter is positive, significant, and of similar magnitude for each of the models where b and c are restricted to be zero (those models where the NIC is symmetric) and is never significant and usually negative[3] when the NIC is allowed either to shift or to rotate (or both). This indicates that it is likely this parameter is capturing some form of asymmetry, unspecified non-linearity, or misspecification in the model rather than a risk premium.

The extended mean equation (Equation 4.8) provides a method of testing for misspecification. Table 4.1 reports likelihood ratio tests for misspecification of the

---

[2] This assumption will be relaxed in Section 4.4.

[3] Of course the fact that it is not significantly different from zero means that its sign cannot be assessed. There are, however, 15 models where b, c, or both are allowed to vary and in 13 cases the sign of γ_1 is negative. The probability of this happening purely by chance is very small.

Table 4.1: LR tests for alternative mean equations

		H_A: $h_t^2, r_{t-1}, r_{t-2}, D_M$	
Model	H_0	LR	p-value
model18 (h_t^2)	$\gamma_2 = \gamma_3 = \gamma_4 = 0$	564.294	< 0.001
model27 (r_{t-1})	$\gamma_1 = \gamma_3 = \gamma_4 = 0$	212.183	< 0.001
model28 (r_{t-1}, r_{t-2})	$\gamma_1 = \gamma_4 = 0$	193.456	< 0.001
model26 (r_{t-1}, D_M)	$\gamma_1 = \gamma_3 = 0$	18.973	< 0.001
model22 (r_{t-1}, r_{t-2}, D_M)	$\gamma_1 = 0$	0.000	1.000

The table presents the results for likelihood ratio tests for the functional form of the conditional mean equation. In each case the test compares a model obtained by restricting one or more parameters of Equation 4.8 to be zero to the model which imposes no restrictions. The likelihood ratio statistics are all distributed χ^2 with between one and three degrees of freedom. The corresponding p-values are included in the final column. A summary of the composition of the models is presented in Table A.1 in Appendix A. The data set comprises excess returns on an U.S. equity index covering the period January 2, 1926 to December 29, 1990, a total of 17,484 observations.

conditional mean equation for the Extended GARCH model selected in Chapter 3 for further development. The table compares the unrestricted conditional mean equation to the restricted equations comprising those with all reasonable subsets of the parameters included.[4] The test results show that the null hypothesis of the simple GARCH-M mean equation ($\gamma_2 = \gamma_3 = \gamma_4 = 0$) can be rejected in favour of the unrestricted model which includes in addition the two lagged returns and the Monday dummy. Similarly, the mean equations which include a single lagged return, both lagged returns alone, and a single lagged return plus the Monday dummy can be rejected in favour of the unrestricted mean equation. The mean equation which includes both lagged returns and the Monday Dummy but excludes the GARCH-M term cannot be rejected in favour of the unrestricted mean equation. Indeed, the difference in the likelihood values does not appear until the third decimal place. This indicates that a conditional mean equation which includes the daily excess return lagged one and two periods and a Monday dummy but not the GARCH-M term is superior to the simple GARCH-M specification.

Finally, Table 4.2 repeats the final panel of Table 3.3 and Table 4.3 repeats the final panel of Table 3.2 substituting the conditional mean specification represented by Equation 4.11 for the specification represented by Equation 3.2. Table 4.2 indicates that the model selected in Section 3.2[5] for further development is still appropriate; the symmetric model (model21, $b = 0$ and $c = 0$) can be rejected in favour of all models which allow some form of asymmetry. The model which allows a rotated but not shifted NIC (model23, $b = 0$ and c free) can also be rejected in favour of a model which allows for both shift and rotation. The model which allows for a shift but no rotation of the NIC (model22, $c = 0$ and b free) cannot be rejected at the 1% level, however. As the conditional variance specification associated with model22 is

---

[4] Table D.1 in Appendix D presents the full set of parameter estimates for these models.

[5] model18 in Table 3.2, which corresponds to model22 in Table 4.3

Table 4.2: LR tests for asymmetry (final mean equation)

	H_A		
H_O	$c = 0, b$ free	$b = 0, c$ free	b and c free
$b = c = 0$	310.461 (< 0.001)	284.963 (< 0.001)	315.919 (< 0.001)
$c = 0, b$ free			5.457 (0.019)
$b = 0, c$ free			30.956 (< 0.001)

The table presents the results for likelihood ratio tests for the form of asymmetry in the conditional volatility. The first line shows the results for likelihood ratio tests of the null hypothesis of no asymmetry ($b = c = 0$) vs. the alternative of asymmetry in some form. The second line shows the results for tests of the null hypothesis of a NIC which is shifted ($c = 0, b$ free) vs. the alternative of a curve which is both shifted and rotated (b, c free), and the third line shows the results for tests of the null hypothesis of a NIC which is rotated ($b = 0, c$ free) vs. the alternative of a curve which is both shifted and rotated. The likelihood ratio statistics are all distributed as χ^2 with either one or two degrees of freedom. The corresponding p-values are included in parentheses. The data set comprises excess returns on an U.S. equity index covering the period January 2, 1926 to December 29, 1990, a total of 17,484 observations.

the Extended GARCH specification selected for further development in Chapter 3, and as the results in Table 4.2 parallel those presented in Table 3.3, development will proceed with the Extended GARCH conditional variance specification and the conditional mean equation:

$$r_t = \gamma_0 + \gamma_2 r_{t-1} + \gamma_3 r_{t-2} + \gamma_4 D_M + u_t \tag{4.11}$$

4.2 Conditional Volatility Equation Development

The previous section provided development and justification for an alternative conditional mean equation. The following section provides similar development and justification for the conditional volatility portion of the model.

The Extended GARCH conditional volatility specification selected at the end of Chapter 3 for further development is given by:

$$h_t = \left(\omega + \alpha h_{t-1}^{\lambda} |v_{t-1} - b|^{\nu} + \beta h_{t-1}^{\lambda}\right)^{\frac{1}{\lambda}} \tag{4.12}$$

The bottom panel of Figure 4.1 shows the standardized residuals from Equation 4.11 where the v_t have been obtained from Equation 4.7 and h_t is as defined in Equation 4.12. If the specification for conditional volatility is correct, v_t should be homoskedastic with unit variance. While comparison to the top panel of the figure indicates that the standardized residuals exhibit less heteroskedasticity than their

Table 4.3: Estimates of model parameters (final mean equation)

Model	γ_0	γ_2	γ_3	γ_4	λ	ν	b	c	ω	α	β	Likelihood
model21	$8.250^{[-4]}$ $(6.382^{[-5]})^*$	0.162 (0.009)*	−0.042 (0.008)*	−0.002 $(1.459^{[-4]})^*$	1.022 (0.328)*	1.570 (0.129)*	0	0	$1.000^{[-4]}$ $(1.596^{[-4]})$	0.066 (0.021)*	0.933 (0.018)*	58937.825
model22	$5.856^{[-4]}$ $(5.780^{[-5]})^*$	0.165 (0.009)*	−0.035 (0.008)*	−0.002 $(1.437^{[-4]})^*$	0.941 (0.265)*	1.538 (0.114)*	0.489 (0.056)*	0	$1.861^{[-4]}$ $(2.364^{[-4]})$	0.059 (0.016)*	0.926 (0.018)*	59093.055
model23	$6.037^{[-4]}$ $(5.889^{[-5]})^*$	0.164 (0.009)*	−0.032 (0.008)*	−0.002 $(1.532^{[-4]})^*$	0.995 (0.303)*	1.353 (0.107)*	0	0.359 (0.051)*	$1.192^{[-4]}$ $(1.742^{[-4]})$	0.063 (0.020)*	0.935 (0.018)*	59080.306
model24	$5.860^{[-4]}$ $(5.804^{[-5]})^*$	0.166 (0.009)*	−0.034 (0.008)*	−0.002 $(1.440^{[-4]})^*$	0.949 (0.266)*	1.493 (0.107)*	0.379 (0.087)*	0.095 (0.062)	$1.738^{[-4]}$ $(2.228^{[-4]})$	0.058 (0.016)*	0.928 (0.018)*	59095.784

Entries of the form $x^{[-y]}$ represent scientific notation; that is $2.5^{[-2]}$ represents the value 2.5×10^{-2}.

The table presents the parameter estimates for nested GARCH models comprising the final conditional mean equation, Equation 4.11 in the text, and parameter restrictions imposed on the nested GARCH conditional volatility specification, Equations 3.3 and 3.4 in the text. Numbers in parentheses under the parameter estimates represent asymptotic standard errors. Parameters which are significantly different from zero at the 5% level are indicated with an asterisk. A summary of the composition of the models is presented in Table A.1 in Appendix A. The data set comprises excess returns on an U.S. equity index covering the period January 2, 1926 to December 29, 1990, a total of 17,484 observations.

Figure 4.1: Residuals from the conditional mean equation

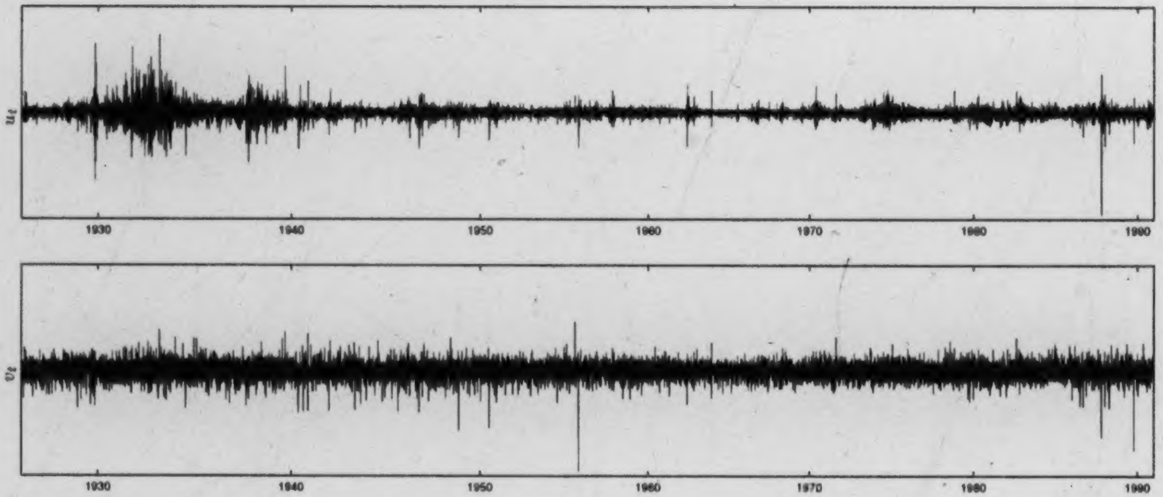

The top panel of the figure presents the residuals from the conditional mean equation given by Equation 4.11. The bottom panel presents the standardized residuals from the same conditional mean equation. The residuals have been standardized as $v_t = u_t/h_t$, where h_t is the conditional volatility defined by Equation 4.12, and should have unit variance. The data set comprises excess returns on an U.S. equity index covering the period January 2, 1926 to December 29, 1990, a total of 17,484 observations.

Table 4.4: Tests of homogeneity of variance

Residual Series	F-statistic	p-value
non-standardized	216.69	< 0.001
standardized	8.37	< 0.001

The table presents the results for a test of homogeneity of variance of the residuals and standardized residuals from Equations 4.11 and Equation 4.12. The test statistic is based on the method of Levene (1960) and is distributed F with 9 and 17470 degrees of freedom. The corresponding p-value is included in the third column. The data set comprises excess returns on an U.S. equity index covering the period January 2, 1926 to December 29, 1990, a total of 17,484 observations.

non-standardized counterparts, a formal test of homogeneity of variance is appropriate to determine if the standardized residuals are indeed homoskedastic. The procedure adopted here is based on the procedure suggested by Levene (1960) as presented in Jobson (1991). This procedure has been chosen over the more common *Bartlett's Test* as the tests performed in Section 3.3 indicate that the standardized residuals are not normally distributed, and Jobson notes that Bartlett's Test is extremely sensitive to non-normality.

To conduct the test the residual series is broken into ten sub-samples each containing 1748 residuals, with the final 4 residuals discarded in the interests of equality of group size. Levene's Method is then applied to the ten groups, and the resulting F-statistics and corresponding p-values presented in Table 4.4. The results in this table indicate that homogeneity of variance can be rejected for both the non-standardized and standardized residuals. Rejection of homogeneity for the non-standardized residuals is expected, as the existence of heteroskedasticity is the feature of the data which the GARCH methodology is intended to model. Rejection of homogeneity of variance for the standardized residuals is problematic, however, as it indicates that the GARCH specification provided in Equation 4.12 does not adequately compensate for the heteroskedasticity apparent in the data. The result is of particular significance as the GARCH specification presented in Equation 4.12 is the Nested GARCH of Hentschel (1995), indicating that the model which best fits this data-set from a large range of members of the GARCH family is unable to remove the heteroskedasticity completely.

While conditional volatility specifications from the GARCH family provide a mechanism for the variance of residuals from the mean equation to change through time, a key assumption underlying conventional GARCH models is that the nature of these changes in the variance is evolutionary; the conditional volatility in the current period is a function of prior values of the conditional volatility and unexpected shocks in prior periods. Because each parameter of the conditional volatility equation is time-invariant, the structure carries with it an implicit assumption that the nature of the determinants of volatility in the market are static. The volatility is determined solely by a combination of news as represented by unexpected shocks from the mean equation and by past values of volatility. There is no allowance for shifts in the underlying unconditional volatility of the market which might be precipitated by structural changes in the economy or by changes in the structure or nature of the market itself. The remaining heteroskedasticity in the standardized residuals demonstrated in Table 4.4, in conjunction with the failure of conventional GARCH models to allow for possible shifts in the unconditional volatility of the market, leads to the development of the Regime Dependent GARCH (RD-GARCH) model. It is this new model which is the focus this study.

4.2.1 Regime Dependent GARCH

The specification for the variance equation can be extended to allow for changes in regimes by defining two $(0, 1)$ dummy variables, D_L and D_H which are zero in periods of normal volatility and one in periods of low or high volatility respectively. These dummy variables are then used to condition each of the parameters ω, α, β, b, λ, and ν as follows:

$$\text{Parameter} = \text{Parameter}_B + \text{Parameter}_L D_L + \text{Parameter}_H D_H$$

that is, each of the parameters has been allowed to vary from the base regime, (B), for both the low, (L), and high, (H), volatility regimes. For example, the λ in Equation 4.12 is replaced by:

$$\lambda = \lambda_B + \lambda_L D_L + \lambda_H D_H \tag{4.13}$$

This substitution allows for a rich parameterization of the model since it allows each estimated parameter to vary across the regimes of interest. It does introduce two complications, however. The first issue that must be considered is a method of determining the dates and extents of the proposed regimes. The second issue is, as noted by Hamilton and Susmel (1994):

> ... The objective is to select a parsimonious representation for the different possible regimes. A specification in which all of the parameters change with each regime would likely be numerically unwieldy and overparameterized ...

The issue of locating proposed regimes is considered next. The question of parameterization is an important issue, and will be addressed in Section 4.2.4. Fortunately, the estimation procedures developed in Chapter 2 are sufficiently robust that the problem of numerical unwieldiness did not materialize.

4.2.2 Regime Detection

In order to estimate models where the parameters are allowed to vary between regimes, it is necessary to define where the regimes begin and end. One possibility would be to search historical accounts manually looking for the dates of events which might have had an impact on the underlying unconditional volatility of the market. This was, in essence, the approach chosen by Baldauf and Santoni (1991) which uses the date of the advent of program trading in a model which seeks to determine if program trading has had a significant impact on the underlying volatility of the market. This procedure is flawed for several reasons: It requires a subjective assessment as to which events should have had an impact, and also makes the implicit assumption that the regime shift would happen at the exact time the news is released; that is, it does not allow for any lead or lag on the part of the market. In the case of Baldauf and Santoni's model, the rejection of the hypothesis of a change in unconditional volatility related to the advent of program trading is, in fact, a rejection of a joint hypothesis regarding a change in unconditional volatility *and* that change happening on the particular date they selected. It does not allow for the possibility that the change in unconditional volatility associated with the advent of program trading either led or lagged the date of the advent of trading on index futures, which was the date selected.

An alternative is to allow the data to indicate where these shifts in volatility take place. This does introduce a risk of data mining, but as long as there is some objective standard for the classification of time periods into regimes, and as long as this standard is determined *a priori* rather than by examination of the data, the risk is minimized.

Several different approaches to locating shifts in market volatility have appeared in the financial literature. Hamilton (1989) suggests a technique, based on a Markov process, which provides probabilities that the economy has switched from one volatility state into another volatility state. This methodology is tied intricately to a particular model specification and, as Hamilton and Susmel (1994) note, cannot in general be applied to the full range of GARCH models because the inclusion of the lagged conditional volatility would make the estimation of the models computationally infeasible, as it would require the maintenance of extremely long Markov chains. Although Gray (1996) suggests an approach which would allow the technique to be applied to a more general class of GARCH models, he provides a solution only for the conventional GARCH(1,1) specification. In addition, his model makes the distributional assumption of conditional normality for disturbances, an assumption which is problematic given the fat-tailed nature of conditional stock returns. The effect of relaxing the conditional normality assumption for the RD-GARCH model is addressed in Section 4.4 for within-sample model fit, and in Section 7.6.4 for out-of-sample forecast performance.

Haugen et al. (1991) present a methodology for detecting variance shifts in which the return series is split into a sequence of data blocks, and comparison of the variance within each adjacent pair of blocks is conducted via an F-test. A significant F-statistic on the test indicates that a difference exists in the variance between the two blocks of data. This procedure provides evidence that the variance has either increased or decreased, but does not directly provide separation of the data into regimes of high, low, and normal volatility as required by the proposed RD-GARCH model; in particular, it is possible for the procedure to suggest consecutive increases or consecutive decreases in variance. The approach suggested by Haugen et al. for dealing with consecutive volatility changes of the same type (either both increases or both decreases) within overlapping data blocks is to consider the both changes to be related to a single event, occurring at the date associated with the largest value for the test statistic. They do, however, allow consecutive volatility changes of the same type provided the changes are sufficiently far apart that they do not occur within overlapping data blocks. In addition, an assumption underlying the F-test is that the return series is independently and normally distributed with zero mean. Unfortunately, significant coefficients estimated for γ_2 and γ_3 in Equation 4.11 indicate violation of the independence assumption, and as is shown in Table 5.2 in Chapter 5, stock returns exhibit significant skewness and excess kurtosis.

The procedure suggested by Tsay (1988) is similar to the procedure of Haugen et al., although somewhat simpler. Tsay's procedure looks at the ratio of the variance of the residual series from an ARIMA model fitted to the data:

$$\phi(B)y_t = \theta(B)u_t \tag{4.14}$$

where the variance ratio, $\hat{R}_d$ is defined as:

$$\hat{R}_d = \frac{d-1}{n-d+1} \times \frac{\sum_{t=d}^{n} u_t^2}{\sum_{t=1}^{d-1} u_t^2} \tag{4.15}$$

where

$$d_r \leq d \leq n - d_r$$

and where d_r is an arbitrarily chosen number of observations reserved at the beginning and end of the series to make it possible to compute the variance ratio. The date of the variance shift, d, is the point at which $\hat{R}_d$ or $\hat{R}_d^{-1}$ is the largest. Since the distribution of $\hat{R}_d$ is unknown, Tsay suggests empirically determined critical values of 3.5, 3.0, or 2.5.

This procedure is not directly applicable to the detection of regimes as defined here for two reasons. By looking in the residuals rather than the data series itself, the proposed exogenous nature of the regimes is ignored. In addition, it only allows for the detection of single points where the variance of a series undergoes a large shift; it does not allow for the detection of multiple regimes of unusually high or low volatility. The first of these concerns is easily dealt with by considering the return series directly. The second is resolved with a new regime detection and isolation procedure.

4.2.3 Expanding Window Variance Ratio Regime Detection

The Expanding Window Variance Ratio (EWVR) regime detection procedure considers a window of observations which moves through the return series. The procedure computes the ratio of the variance of the observations inside the window to the variance of the observations outside the window and sequentially isolates and excludes the largest window of observations with a variance ratio greater than a given critical value. More formally:

$$\hat{R}_d = \frac{n-w-1}{w-1} \times \frac{\sum_{t=d}^{d+w}\left(r_t - \bar{r}_{iw}\right)^2}{\sum_{t=1}^{d-1}\left(r_t - \bar{r}_{ow}\right)^2 + \sum_{t=d+w+1}^{n}\left(r_t - \bar{r}_{ow}\right)^2} \qquad (4.16)$$

where

$$1 \leq d \leq n$$

and where $\bar{r}_{iw}$ is the mean return inside the window, $\bar{r}_{ow}$ the mean return outside the window, and w the size of the window, which is selected based on criteria outlined below. Here d is allowed to vary $1 \leq d \leq n$ because there is no need to reserve observations at the beginning and end of the series to compute a variance ratio. The beginning date of the regime, d, is once again the point at which $\hat{R}_d$ or $\hat{R}_d^{-1}$ is the largest. Once a likely regime has been detected based on the critical value of 3.0 (chosen because it is the middle of the range suggested by Tsay), an iterative procedure is adopted to determine the extent of the regime:

1. An additional observation is added onto the beginning of the window and the variance ratio recomputed.
2. The additional observation which was added onto the beginning of the window is removed, an additional observation added onto the end of the window, and the variance ratio recomputed.
3. If both steps 1 and 2 caused the variance ratio to drop below the critical value, iteration is stopped and the current window is declared to be a regime.
4. Otherwise, the observation from step 1 or 2 which made the smallest change to the variance ratio is retained and iteration continues with step 1.

Once the full extent of the regime has been detected in this manner, the observations within the window are excluded from the sample and the overall procedure is repeated until no more regimes are found.[6]

The criteria for selecting w, the window size, are necessarily ad-hoc and based on the proposed nature of and causes for the regimes. If the cause for the suggested shifts in the unconditional variance of the return series is shifts in the nature of the economy, a larger value of w would be expected as such regimes should be long-lived. If, on the other hand, the cause for the apparent shifts is related to internal dynamics of the market, a smaller value for w would be appropriate. A value of 60 trading days has been chosen as the starting window size, approximately representing a calendar quarter. The window can increase in size as observations are added, but cannot shrink. This value should be small enough to detect even relatively short-lived economic events which might be considered regimes, as well as allowing for some forms of market dynamics. It should also be large enough to avoid detecting very short fluctuations, which can be considered to be noise.

---

[6] See Appendix B for an illustrative example.

4.2.4 Correlation Among Explanatory Variables

The rich parameterization of the variance equation provided by the regime dependent framework leads to a problem of correlation among the explanatory variables. Since each of the six explanatory variables in the variance equation is conditioned by the same pair of $(0,1)$ dummies, the 12 regime specific variables appear to be in two highly correlated groups. This has the effect of greatly overstating the standard errors reported for these variables, making standard t-tests for significance of the coefficients invalid. As an alternative to using the asymptotic standard errors calculated from the ML estimation, a simple stepwise procedure was adopted which is analogous to the backward elimination method for the multiple linear regression model:

1. The model is estimated with all 18 conditional volatility equation parameters included.
2. Each of the 12 sub-models created by removing one of the parameters conditioned by a regime dummy is also estimated.
3. The sub-model which made the smallest change to the likelihood value (the one with the smallest likelihood ratio) is chosen as the new unrestricted model and the procedure returns to step 2 with the 11 sub-models of the new unrestricted model.
4. Iteration continues until the likelihood ratio for removal of any of the remaining parameters is significant at the 1% level.
5. The model which results is then reported as the regime dependent model.

4.3 The Final Model

The mean equation for the final model is defined by Equation 4.11, and the conditional variance equation is defined by Equation 4.17 with the regime dummy variables included for each of the estimated parameters in the variance equation:[7]

$$\frac{h_t^{\boldsymbol{\lambda}} - 1}{\boldsymbol{\lambda}} = \boldsymbol{\omega}' + \boldsymbol{\alpha}' h_{t-1}^{\boldsymbol{\lambda}} |v_{t-1} + \boldsymbol{b}|^{\boldsymbol{\nu}} + \boldsymbol{\beta} \frac{h_{t-1}^{\boldsymbol{\lambda}} - 1}{\boldsymbol{\lambda}} \tag{4.17}$$

Together these equations define the Regime Dependent GARCH (RD-GARCH) specification. The model provides a total of 10 parameters in the conventional form and, as the six conditional volatility parameters ω, α, β, b, λ, and ν each become three parameters when conditioned by the regime dependent dummy variables, a total of 22 parameters in the regime dependent form. Because the specification of Equation 4.17 is recursive in nature and requires a value for h_{t-1} to compute a value for h_t it necessary to provide a value for h_0, the initial conditional volatility. Two different methods were investigated for specifying a value for h_0: estimation as a nuisance parameter, and using an estimate of the square root of the unconditional variance. Both methods gave nearly identical results. The optimization routine tended to converge somewhat faster when h_0 was estimated as a nuisance parameter so this method has been adopted, making a total of 11 parameters in the conventional form and 23 in the regime dependent form.

---

[7] Parameters which are conditioned by the regime dummy variables are indicated in bold.

4.4 Alternative Distributional Assumptions

Table 3.5 presented additional descriptive statistics for the standardized residuals from the mean equation for the models estimated in Chapter 3. As outlined in Section 3.3, the skewness and kurtosis measures presented in that table, in conjunction with the Jarque-Bera test statistics which were also presented in the table, indicate that the assumption of normality under which the likelihood equation was derived is problematic.

This section considers two alternative distributions for the disturbances: the generalized error distribution and the Student-t distribution, both of which allow the kurtosis of the disturbances to be adjusted, and both of which nest the standard normal distribution as a special case. Because these distributions both nest the normal, likelihood ratio tests can be performed to determine if the alternative distributions provide superior model fit to the data. Attempts were made using the Pearson Curve approach suggested by Freund (1962) to derive a density function which nests the normal distribution and allows both the skewness and kurtosis to be adjusted, however the expression which resulted was numerically intractable.

4.4.1 The Generalized Error Distribution

For the generalized error distribution (GED), the density of v_t, where v_t has zero mean and unit variance, is:

$$f(v_t) = \frac{k \exp\left[-\frac{1}{2}\left|\frac{v_t}{\Lambda}\right|^k\right]}{\Lambda 2^{(1+\frac{1}{k})}\Gamma(1/k)} \tag{4.18}$$

where

$$\Lambda \equiv \left[2^{-\frac{2}{k}}\frac{\Gamma(1/k)}{\Gamma(3/k)}\right]^{\frac{1}{2}} \tag{4.19}$$

and

$$-\infty < v_t < \infty$$
$$0 < k \leq \infty$$

The mean and variance have been standardized to zero and one respectively, and k is a parameter which controls the thickness of the tails. Smaller values for k are associated with thicker tails. Values of $k > 2$ correspond to platykurtic distributions, and values of $k < 2$ correspond to leptokurtic distributions. For the case where $k = 2$, the density of the GED is that of the standard normal distribution. This parameter may be allowed to vary across regimes in the same manner as the other parameters of the conditional volatility equation. In this way, the kurtosis is free to change across the various regime types. Derivation of the log-likelihood equation for the GED is provided in Appendix C.2.

4.4.2 The Student-t Distribution

For the Student-t distribution, the density of v_t is given by:

$$f(v_t) = \frac{\Gamma[(k+1)/2]}{\sqrt{\pi k}\Gamma[k/2]}\left(1+\frac{v_t^2}{k}\right)^{\frac{-(k+1)}{2}} \tag{4.20}$$

where k represents the number of degrees of freedom, and in a manner similar to the GED controls the thickness of the tails. Values of $1/k > 0$ correspond to leptokurtic distributions and larger values of $1/k$ are associated with thicker tails. The density has zero mean but variance given by:

$$\operatorname{var}(v_t) = \frac{k}{k-2} \tag{4.21}$$

which, for large values of k, would approach the required property of unit variance. Unfortunately, as excess kurtosis is associated with smaller values for k, this density function is not appropriate. The standardized Student-t distribution has the required property of zero mean and unit variance. For the standardized Student-t distribution, the density of v_t is given by:

$$f(v_t) = \frac{\Gamma[(k+1)/2]}{\sqrt{\pi(k-2)}\Gamma[k/2]}\left(1+\frac{v_t^2}{k-2}\right)^{\frac{-(k+1)}{2}} \tag{4.22}$$

As $k \to \infty$, the density of both the standardized and non-standardized Student-t distributions approach that of the standard normal distribution. As with the GED, this parameter may be allowed to vary across regimes and thus to allow the kurtosis to change across the various regime types. Derivation of the log-likelihood equation for the Student-t distribution is provided in Appendix C.3.

Chapter 5

Application of RD-GARCH

This chapter takes the RD-GARCH model developed in Chapter 4 and subjects it to a number of tests to determine how well it performs at describing volatility of the U.S. equity index over the period January 2, 1926 to December 29, 1990. The chapter considers five topics related to the application of the RD-GARCH methodology: the results from the regime detection procedure is discussed first, with an analysis of the regimes detected. The behaviour of the statistical moments of the return series during the three types of volatility regimes is then considered, followed by a discussion of the overall performance of the RD-GARCH methodology relative to conventional GARCH models. An in-depth analysis of the performance of the particular RD-GARCH model selected in Section 4.3 is then conducted, both under the assumption of normally distributed disturbances and under the alternative distributional assumptions outlined in Section 4.4.

5.1 Regime Detection

Application of the procedures outlined in Section 4.2.3 to the excess returns dataset described in Section 3.1.1 resulted in the isolation of the regimes identified in Table 5.1. In addition to the beginning and ending dates identified for each regime, the mean and unconditional variance of the excess return series are also reported. Figure 5.1 shows the excess return series with regimes of high volatility indicated by boxes above the zero axis, and regimes of low volatility indicated by boxes below the zero axis.

It would be useful to associate economic events with the detected regimes. Of course such association after the fact is purely descriptive, although it does provide a "sanity check" on the regime detection process.

It is convenient to deal with the high volatility regimes first. Inspection of Figure 5.1 indicates that it would be reasonable to consider the eight separate high volatility regimes to comprise five groups. The three regimes covering the periods March 21, 1929 through January 17, 1939, March 14 through December 11, 1939, and February 2 through October 10, 1940 can quite naturally be considered together. This period covers the events leading up to the market crash of October 1929, and the

Table 5.1: Dates of regimes

Base		Low Volatility				High Volatility			
$\bar{r}$	Var	From	To	$\bar{r}$	Var	From	To	$\bar{r}$	Var
3.531×10^{-4}	6.089×10^{-5}	1926 11 09	1927 01 26	2.434×10^{-5}	2.194×10^{-5}	1929 03 21	1939 01 17	-1.103×10^{-4}	4.052×10^{-4}
		1927 03 23	1927 06 04	1.186×10^{-3}	2.218×10^{-5}	1939 03 14	1939 12 11	-3.583×10^{-5}	2.281×10^{-4}
		1944 01 06	1944 03 24	3.503×10^{-4}	2.117×10^{-5}	1940 02 23	1940 10 07	-5.290×10^{-4}	2.037×10^{-4}
		1944 04 19	1944 07 19	1.490×10^{-3}	2.098×10^{-5}	1946 07 09	1947 02 17	-6.991×10^{-4}	1.977×10^{-4}
		1950 01 13	1950 03 31	6.748×10^{-4}	2.015×10^{-5}	1973 10 22	1974 02 06	-2.411×10^{-3}	1.806×10^{-4}
		1951 07 25	1951 10 20	7.833×10^{-4}	2.365×10^{-5}	1974 05 22	1975 05 08	7.461×10^{-5}	1.884×10^{-4}
		1951 11 08	1952 01 31	1.269×10^{-3}	2.084×10^{-5}	1982 07 20	1982 11 12	3.154×10^{-3}	1.775×10^{-4}
		1952 02 20	1952 10 16	3.062×10^{-4}	2.186×10^{-5}	1986 08 25	1988 04 28	7.283×10^{-5}	2.617×10^{-4}
		1952 11 03	1953 02 09	1.006×10^{-3}	2.202×10^{-5}				
		1953 10 30	1954 04 30	1.322×10^{-3}	2.096×10^{-5}				
		1957 02 14	1957 07 30	1.012×10^{-3}	2.038×10^{-5}				
		1961 01 05	1961 04 04	1.981×10^{-3}	2.034×10^{-5}				
		1961 09 26	1962 04 05	3.625×10^{-4}	1.940×10^{-5}				
		1962 05 29	1970 06 01	1.828×10^{-4}	4.232×10^{-5}				
		1971 01 28	1971 05 18	8.945×10^{-4}	2.235×10^{-5}				
		1972 01 25	1972 04 26	8.345×10^{-4}	2.059×10^{-5}				
9860 Days		3279 Days				4345 Days			

The table presents the dates, mean excess return, and unconditional variance of the excess return for each of the regimes identified. The base regime covers the entire sample period from January 2, 1926 through December 29, 1990 (17,484 observations), with the returns for the high and low volatility regimes excluded from the calculation of the sample statistics. Each row in the sections of the table labeled "Low Volatility" and "High Volatility" represent a single period identified as either a regime of unusually low or unusually high unconditional volatility respectively. The number of days identified at the end of each column identifies the total number of days identified with that particular type of regime in the sample period.

Figure 5.1: Excess returns 1926–1990

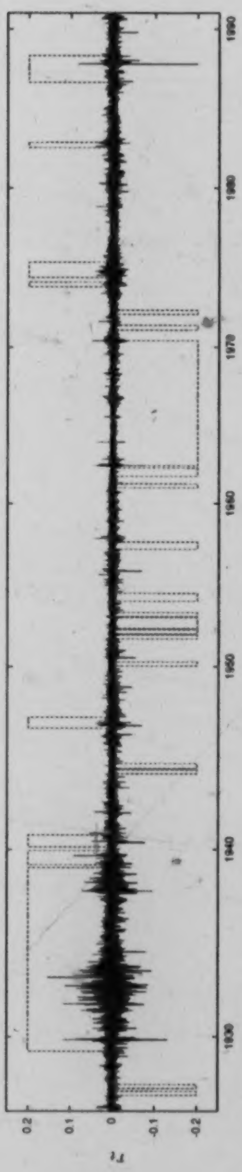

The figure presents the excess returns on the stock index series for the period January 2, 1926 through December 29, 1990 (17,484 observations). Regimes of unusually high or low volatility in the returns are indicated by boxes above or below the zero axis respectively.

subsequent Great Depression. The period just prior to the crash was characterized by significant margin speculation in the U.S. markets, causing the Federal Reserve Board to issue a directive to its member banks on February 2, 1929 forbidding the granting of margin loans. The Depression following the Crash was also characterized by extremely high volatility, an effect which lasted until just over one year before the United States entered World War II. Similarly, the two regimes covering the periods October 22, 1973 through February 2, 1974 and May 22, 1974 through May 8, 1975 are close enough together to be considered to be part of the same economic event. This period corresponds to the Arab oil embargo and subsequent oil crisis in the United States, as well as the resignation of President Nixon on August 8, 1974. The U.S. economy underwent a major recession during this time, the largest since the depression of the 1930's. The high volatility regime covering the period August 25, 1986 through April 28, 1988 captures the market crash of October 1987. The other two high volatility regimes are more difficult to explain. The regime covering the period July 9, 1946 through February 17, 1947 occurs slightly less than a year after the end of the Second World War, and coincides both in starting date and duration with the announcement on July 15, 1946 that the war-time wage and price controls were to be extended for an additional year. Why this would cause an increase in the volatility in the stock market is unclear at best and may simply be a coincidence. The final regime, July 20 through November 12, 1982 is even more difficult to interpret. It is the shortest of the independent high volatility regimes, and also has the largest mean return of any high *or* low volatility regime. This is particularly strange as the mean return in all of the other high volatility regimes is lower than the mean return in either the base or low volatility regimes. This could simply be a result of sampling variation.

It is more difficult to provide interpretation for the periods of low volatility. Events such as depressions and recessions which seem to have a destabilizing impact on the market are relatively well defined and documented, while events which reduce the volatility seem to be difficult to isolate. One possible reason for this might be the nature of the impact of different forces on the market. An event which has a destabilizing impact, such as a market crash, tends by its very nature to have a sudden onset. On the other hand, more stable market conditions could reflect a lack of events. Figure 5.1 indicates that the 16 separate low volatility regimes can be considered to belong to six general groups. The two regimes covering the periods November 9, 1926 through January 26, 1927 and March 23 through June 4, 1927 belong to one such group. The second group comprises the regimes covering the periods January 6 through March 24, 1944 and April 19 through July 19, 1944. The third group covers the five separate regimes which span the period January 1, 1961 through April 26, 1972. The remaining seven low volatility regimes are not nearly as tightly grouped. It seems reasonable, however, to consider the five regimes covering the period July 25, 1951 through April 30, 1954 to be a group, and the regimes covering the periods January 13 through March 31, 1950 and February 14 through July 30, 1957 to stand on their own. Other than the observation that the detected regimes tend to fall into these groups, and the additional observation that three of the six groups correspond to the involvement of the United States in wars (World War II, Korea, and Vietnam), obvious tie-ins to events in the economy are lacking.

Table 5.2: Descriptive statistics for regimes

Regime	Mean	Variance	Skewness	Kurtosis
Overall	2.577×10^{-4}	1.264×10^{-4}	−0.357	20.007
Base	3.526×10^{-4}	6.089×10^{-5}	−0.467	4.089
Low	4.463×10^{-4}	3.428×10^{-5}	0.079	6.482
High	-9.998×10^{-5}	3.445×10^{-4}	−0.194	8.879

The table presents descriptive statistics for the excess return series used in this section. The first row presents the statistics for the entire sample period, with no adjustments made for regimes of differing volatility. The second row presents the statistics for the entire sample period with the returns for periods identified as having either unusually high or unusually low volatility removed from the sample. The third and fourth rows present the statistics for all of the periods identified as either unusually high volatility or unusually low volatility combined by type of regime. The data set comprises excess returns on an U.S. equity index covering the period January 2, 1926 to December 29, 1990, a total of 17,484 observations.

5.2 Behaviour of Returns During Regimes

Table 5.2 reports descriptive statistics of the excess return series both for the entire sample period, and summarized by the type of regime. By construction, the variance is highest in the high volatility regimes, and lowest in the low volatility regimes. The critical value for the variance ratio test which was used to separate the observations into the appropriate regimes was 3.0, so it would be expected that the variance in the high and low regimes should differ from the variance in the base regime by approximately a factor of three. In fact, the variance in the high volatility regime is 5.7 times larger than the variance in the base regime, and the variance in the low volatility regime is only 1.8 times smaller than the variance in the base regime. This difference is due to the order in which the regimes were detected. The first regime to be detected and removed was the low volatility period from May 29, 1962 through June 1, 1970. Removing this low volatility regime caused the volatility of the observations remaining in the base regime to increase. The second regime detected and removed was the high volatility period which covered the period March 21 1929 through May 26, 1930. Removal of this high volatility regime caused a decrease in the volatility of the observations remaining in the base regime. The removal of observations causing a shift in the volatility of the "base" regime produces the effect of an apparent violation of the variance ratio in the final distribution of observations into the three types of regimes. The impact of this effect will be investigated further in Chapter 6 where the stability of the RD-GARCH methodology is addressed.

The mean return in high volatility regimes is lower than in either the low or the base regimes, and is in fact negative. Although the difference is not statistically significant, examination of Table 5.1 shows that eight high volatility regimes have been detected, and that the mean return is lower than the return in the base regime in seven of the cases. Similarly, the mean return in the low volatility regimes is higher than the mean return in either the base or high volatility regimes. Again, this difference is not statistically significant, but a total of 16 low volatility regimes have been detected and in 12 cases the mean return is higher than the mean return in the base regime. In 15 cases, the mean return in the low volatility regimes is higher than the overall mean return in high volatility regimes. This would tend to indicate that periods of high volatility are characterized by "bear" markets, and that periods of low volatility are characterized by "bull" markets.

The final two columns in Table 5.2 report skewness and excess kurtosis statistics. In keeping with results commonly found in stock returns, the returns are negatively skewed overall, and are also negatively skewed in both the base and high volatility regimes. They are slightly positively skewed in low volatility regimes, however. The excess kurtosis that is commonly found in stock returns is present, but when the returns are separated into periods of high, medium, and low volatility the level of excess kurtosis declines in each type of regime. These findings indicate that the distribution of stock returns is not time-invariant, and provide additional incentive for the tests of alternative distributions conducted in Section 5.5.

5.3 Overall Performance of RD-GARCH

As an initial test, the Regime Dependent GARCH methodology was applied to the 20 models estimated in Chapter 3. The resulting parameter estimates and diagnostic statistics are presented in Tables D.3 and D.4 respectively in Appendix D. Results of likelihood ratio tests comparing each of the 20 regime dependent models to their conventional counterparts are presented in Table 5.3 and indicate that, without exception, the regime dependent model is superior to the conventional model. Likelihood ratio statistics are in the range 373–606, and are distributed χ^2 with up to 12 degrees of freedom, indicating extremely strong rejection of the null hypothesis of no difference between the restricted (conventional) and unrestricted (regime dependent) models. This indicates that the Regime Dependent GARCH methodology results in a model fit which is superior to that obtained by an otherwise identical model which does not allow the GARCH parameters to vary across volatility regimes.

5.4 Further Analysis of the Selected Model

The parameter estimates from the model represented by Equations 4.11 and 4.17 in both conventional and regime dependent forms are presented in Table 5.4. In the table, the first section reports the results for the conventional version of the model—the model which imposes the restriction that all parameters conditioned by regime dependent dummies are zero. The second section reports the results for the regime dependent version of the model—the model which does not impose restrictions on the regime dependent dummies, and uses the stepwise procedure outlined in Section 4.2.4 to select the best model. Variables which have been excluded from the final model by the stepwise procedure are reported with zero for the parameter estimate and no corresponding standard error.

Comparison of the mean equation parameters, γ_0, γ_2, γ_3, and γ_4, between the two sections of Table 5.4 indicates that almost no difference in the conditional mean

Table 5.3: LR tests for RD-GARCH

Model	LR	p-value
model01	372.993	< 0.001
model02	559.426	< 0.001
model03	526.471	< 0.001
model04	591.121	< 0.001
model05	381.276	< 0.001
model06	584.985	< 0.001
model07	542.179	< 0.001
model08	606.049	< 0.001
model09	407.917	< 0.001
model10	585.948	< 0.001
model11	529.150	< 0.001
model12	587.089	< 0.001
model13	395.281	< 0.001
model14	570.288	< 0.001
model15	517.733	< 0.001
model16	568.305	< 0.001
model17	381.655	< 0.001
model18	595.103	< 0.001
model19	543.741	< 0.001
model20	593.745	< 0.001

The table presents the results for likelihood ratio tests for the inclusion of regime dependence in the GARCH specification. Each row of the table tests the restrictions that, for the corresponding models from Tables 3.2 and D.3, all of the coefficients of the parameters conditioned by the regime dependent dummies, D_L and D_H in Equation 4.17 are zero. The likelihood ratio test statistics are all distributed as χ^2 with between 2 and 12 degrees of freedom. The corresponding p-values are shown in the final column. A summary of the composition of the models is presented in Table A.1 in Appendix A. The data set comprises excess returns on an U.S. equity index covering the period January 2, 1926 to December 29, 1990, a total of 17,484 observations.

results from allowing the conditional variance to depend upon the regime. In both cases, the impact of the prior period's return is positive and significant with a value of approximately 0.16, the return two periods ago has a small negative impact, and the "Monday" effect is very small and negative. This indicates that, for daily data, there is a small but statistically significant positive relationship between the current return and the return one period ago; the prior period return has predictive power for the current return. There is also a significant relationship between the return in the current period and the return two periods ago. This relationship is very small in magnitude and of the opposite sign to the relationship between the current and immediately prior period returns, possibly indicating a slight systematic over-reaction of the current return to the impact of the prior period return. The small negative coefficient of the Monday dummy indicates that Mondays are characterized by slightly smaller than the average returns, a finding which is supportive of the weekend effect.

The conditional variance parameters show considerably more activity than the mean equation parameters. The parameter λ, which controls the Box-Cox transfor-

Table 5.4: Estimates of model parameters (final model)

Regime	γ_0	γ_2	γ_3	γ_4	λ	ν	b	ω	α	β	Likelihood
none detected	$5.856^{[-4]}$	0.165	−0.035	−0.002	0.941	1.538	0.489	$1.861^{[-4]}$	0.059	0.926	59093.055
	$(5.780^{[-5]})^*$	$(0.009)^*$	$(0.008)^*$	$(1.437^{[-4]})^*$	$(0.265)^*$	$(0.114)^*$	$(0.056)^*$	$(2.364^{[-4]})$	$(0.016)^*$	$(0.018)^*$	
base regime	$6.311^{[-4]}$	0.162	−0.034	−0.002	4.721	2.587	0.754	$5.890^{[-12]}$	0.096	0.741	59367.438
	$(5.597^{[-5]})^*$	$(0.009)^*$	$(0.008)^*$	$(1.405^{[-4]})^*$	$(1.652)^*$	$(0.361)^*$	$(0.051)^*$	$(4.785^{[-11]})$	$(0.031)^*$	$(0.073)^*$	
low volatility					−3.794	−1.333	0	$4.407^{[-4]}$	0	0.094	
					$(1.527)^*$	$(0.402)^*$		$(8.079^{[-4]})$		(0.052)	
high volatility					−3.420	−0.905	0	$1.307^{[-4]}$	0	0.106	
					$(1.484)^*$	$(0.442)^*$		$(3.327^{[-4]})$		$(0.050)^*$	

Entries of the form $x^{[-y]}$ represent scientific notation; that is $2.5^{[-2]}$ represents the value 2.5×10^{-2}.

The table presents the parameter estimates for the GARCH model selected in Chapter 4. Numbers in parentheses under the parameter estimates represent asymptotic standard errors. Parameters which are significantly different from zero at the 5% level are indicated with an asterisk. The first two rows of the table are the parameter estimates and corresponding standard errors that result when no regime detection and separation is performed. The remaining six rows present the parameter estimates and standard errors which result when regimes of unusually high and low volatility are detected and the parameters of the GARCH model are allowed to change in the different regimes. Parameters which are excluded by the stepwise estimation procedure are represented by a value of zero, with no corresponding standard error presented. The data set comprises excess returns on an U.S. equity index covering the period January 2, 1926 to December 29, 1990, a total of 17,484 observations.

mation of the conditional volatility series, has the effect of transforming the functional form of the conditional volatility. The purpose of this transformation is to render the conditional volatility both more nearly homoskedastic and more nearly normally distributed, although as Judge et al. (1985) suggest, when this parameter is estimated with the model, the parameter estimate tends to be biased in the direction of transforming the series to be more nearly homoskedastic. In the conventional version of the model, λ is nearly one, indicating that a different functional form provided by the Box-Cox transformation is not effective in improving the normality or homoskedasticity of the conditional volatility. When λ is allowed to vary across regimes, it takes on a value of 4.721 in the base regime. This indicates that evolution through the transformed variance has the effect of reducing either the heteroskedasticity or the non-normality (or both) of the conditional volatility. In both the high and low volatility regimes its value returns to nearly one again (recall that the regime estimates must be added to the base estimates to obtain regime values). Either the separation of the series into regimes has removed the problems of non-normality and heteroskedasticity for the periods of unusually high or low volatility, or the nature of the non-normality and heteroskedasticity is different in these regimes and the Box-Cox transformation is not effective in rendering the series more nearly homoskedastic and normally distributed.

The conditional volatility specification for the conventional GARCH(1,1) model of Bollerslev (1986) is given by:

$$\begin{aligned} h_t^2 &= \omega + \alpha h_{t-1}^2 v_{t-1}^2 + \beta h_{t-1}^2 && (5.1a) \\ &= \omega + \alpha h_{t-1}^2 \left(v_{t-1}^2 - 1\right) + (\alpha + \beta)\, h_{t-1}^2 && (5.1b) \end{aligned}$$

An assumption underlying this model formulation is that $v_t \sim iid(0,1)$. The term $\left(v_{t-1}^2 - 1\right)$ has zero mean regardless of the distribution of v_t, and can be thought of as a volatility shock to the model. In the model, the parameter α measures the impact of this volatility shock on subsequent conditional volatility. The term $(\alpha + \beta)$ measures the impact of the prior period conditional volatility on the current conditional volatility, and represents a measure of the persistence or long-term memory of the conditional volatility. For the Nested GARCH specification used in this work, the relationship between the model parameters and measures of the impact of volatility shocks and persistence are not nearly as clear. The Extended GARCH conditional volatility specification is given by:

$$h_t = \left(\omega + \alpha h_{t-1}^{\lambda} |v_{t-1} - b|^{\nu} + \beta h_{t-1}^{\lambda}\right)^{\frac{1}{\lambda}} \qquad (5.2)$$

The term $|v_{t-1} - b|^{\nu}$ presents a problem in this specification. In the absence of a stronger distributional assumption than $v_t \sim iid(0,1)$, an expression for the expected value of $|v_{t-1} - b|^{\nu}$ cannot be determined.[1] As there is no term analogous to the expression $\left(v_{t-1}^2 - 1\right)$ from Equation 5.1b in Equation 5.2, there is no simple transformation of Equation 5.2 which will provide a specification that separates the impact

---

[1] In fact, even under the assumption of normality for the residuals, the expected value expression is intractable. See Footnote 4 on page 128.

Figure 5.2: News Impact Curve 1926–1990

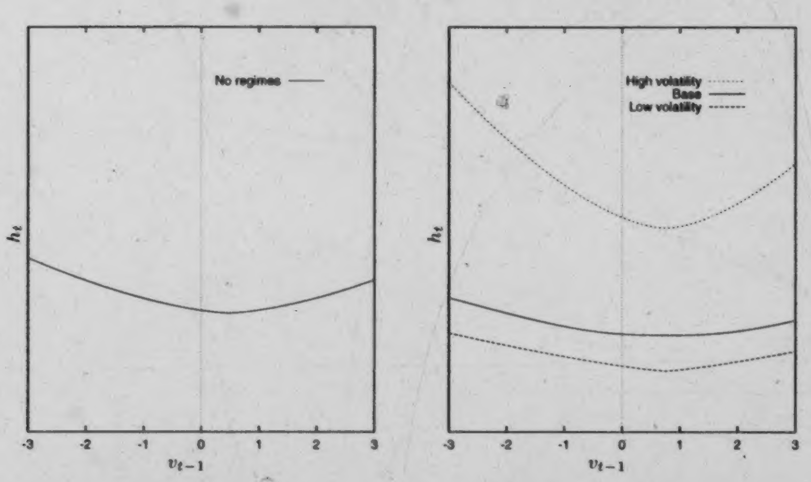

The figure presents the NIC relating volatility shocks in the previous period to the conditional volatility in the current period. The left-hand panel shows the NIC for the conventional version of the model, and the right-hand panel shows the three curves which result when the regime dependent methodology is adopted. The range of volatility shocks over which the figure is presented corresponds to the range ±3. Since $v_t \sim iid(0,1)$ this corresponds to ±3 standard deviations. Both figures include a dashed vertical line corresponding to $v_{t-1} = 0$ for comparison of asymmetry effects. To facilitate comparison, the vertical scale of both figures are the same. The data set comprises excess returns on an U.S. equity index covering the period January 2, 1926 to December 29, 1990, a total of 17,484 observations.

of volatility shocks from the effect of prior conditional volatility. In the absence of a simple functional form for expressing the impact of volatility shocks and persistence, the NIC and HIC respectively will be employed to perform a similar function.

Figure 5.2 presents the NICs for the model under consideration. The left-hand panel shows the curve which results from the conventional version of the model, and the right hand panel shows the three curves which result when the parameters conditioned by the regime dependent dummy variables are included. The range of volatility shocks over which the figure is presented corresponds to the range ±3. Since $v_t \sim iid(0,1)$ this corresponds to ±3 standard deviations. Both figures include a dashed vertical line corresponding to $v_{t-1} = 0$ for comparison of asymmetry effects. To facilitate comparison, the vertical scale of both figures are the same.

In the NIC for the model in which regime detection and isolation was not performed, the single curve is shifted somewhat to the right, and exhibits a small amount of convexity. The shift of the curve is indicative of asymmetry in the response of conditional volatility to the volatility shock in the previous period. Because the curve is shifted to the right a shock of $v_{t-1} = 1$ (one standard deviation), for example, will have a smaller impact on h_t than a shock of $v_{t-1} = -1$ (a negative shock of the same magnitude). If the volatility shocks are viewed as a response to news, this indicates that a positive news event has a smaller effect on subsequent conditional volatility than a negative news event of the same magnitude. Bad news has a larger impact on subsequent conditional volatility than good news. As the NIC exhibits convexity, the response of the conditional volatility to volatility shocks of the same sign but different magnitudes is not linear. The difference in impact between a shock of $v_{t-1} = 2$ and a shock of $v_{t-1} = 3$ is greater than the difference in impact between a shock of $v_{t-1} = 1$ and a shock of $v_{t-1} = 2$. Again, if the volatility shocks are interpreted as a response to news, this indicates that larger news events of either sign have an increasingly large impact on subsequent conditional volatility.

The separation of the curves in the regime dependent model reflects the difference in unconditional volatility found in the three types of regimes. By construction, the low volatility regimes exhibit the smallest level of unconditional volatility, the base regime a larger level, and the high volatility regimes the largest level.

The convexity exhibited by the NIC differs markedly across the three different regimes. High volatility regimes are characterized by a NIC which has a much higher degree of curvature than the curve for either the base regime or low volatility regimes. Conditional volatility is more responsive to large shocks in periods of high volatility, and this effect is not simply due to the higher unconditional volatility that exists in these regimes; the difference in impact between small and large shocks is greater in periods of high volatility than in periods of normal or low volatility. This indicates that in periods of high volatility, large news events, both positive and negative, have a larger impact on the volatility in the market in subsequent periods than similar news events would have during periods of lower volatility. If the volatility in the market is already unusually high, investors tend to react to large news events more aggressively than they would otherwise in periods of lower volatility. Low volatility regimes exhibit a slightly sharper impact curve than the base regime, although it exhibits less curvature. For a certain range of magnitudes of shocks, the difference in impact for slightly larger shocks in low volatility periods is greater than for normal volatility periods. Once outside that range, however, the reverse is true: the difference in impact for much larger shocks is greater in normal volatility periods than in low volatility periods. This difference is quite small, however, and seems to largely mimic the effect noted in the high volatility case; in the base volatility regime investors tend to react more aggressively to large news events than they do in low volatility regimes, although this difference is quite small. The curvature of the NIC for the conventional version of the model appears largely to be an average of the curvature of the three NICs from the regime dependent model. In this case, allowing the model parameters to vary across regimes has provided a mechanism for identifying a difference in market reaction to news items of different magnitude.

The degree by which the curves are shifted to the right is similar for all three curves in the regime dependent model, and is somewhat larger than the shift which is apparent in the conventional version. This indicates that the asymmetry which is apparent in the conventional version of the model is also present, and is in fact more pronounced, when the model parameters are allowed to vary across regimes. It is somewhat surprising that the asymmetry appears increased in all three regime types in the regime dependent model as it might be expected that the degree of asymmetry in the conventional model should be an average of the asymmetry in each of the regimes from the regime dependent model. One possible explanation is

Figure 5.3: History Impact Curve 1926–1990

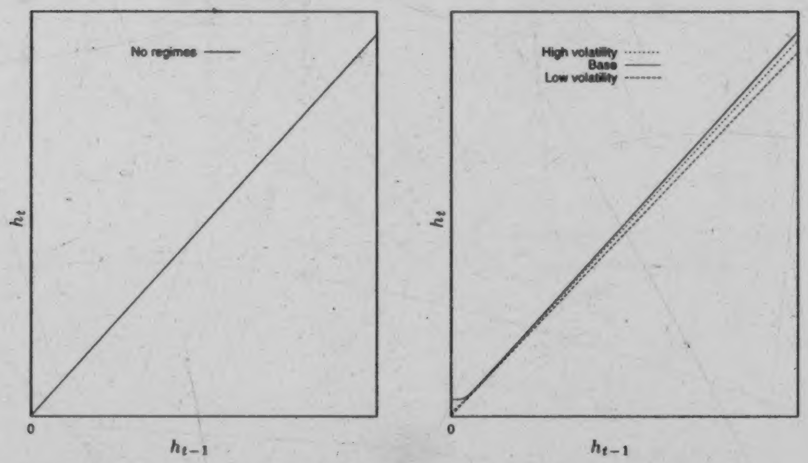

The figure presents the HIC relating the conditional volatility in the previous period to the conditional volatility in the current period. The left-hand panel shows the HIC for the conventional version of the model, and the right-hand panel shows the three curves which result when the regime dependent methodology is adopted. The data set comprises excess returns on an U.S. equity index covering the period January 2, 1926 to December 29, 1990, a total of 17,484 observations.

that interaction between the parameters that control shifting and curvature in the conventional model has caused the curve to be shifted less than appropriate, and that the additional flexibility provided by the regime dependent model has allowed the three NICs to reflect the appropriate level of asymmetry.

Figure 5.3 presents the HICs for the model under consideration. The left-hand panel shows the curve which results from the conventional version of the model, and the right hand panel shows the three curves which result when the parameters conditioned by the regime dependent dummy variables are included. Comparison of the two figures reveals that the large decrease in persistence suggested by the change in value of β over-states the true change in persistence resulting from regime dependence. The level of persistence in the base regime is similar to the overall level of persistence observed from the model which does not allow for regime dependence. There does appear to be a small reduction in the level of persistence in high and low volatility regimes, with the low volatility regimes exhibiting the least persistence in conditional volatility and the high volatility regimes exhibiting slightly more persistence.

Table 5.5 presents the additional diagnostic statistics outlined in Chapter 3. The first section of the table repeats the results from the Extended GARCH model with GARCH-M mean equation (Extended GARCH-M model) selected at the end of Chapter 3 for comparative purposes. The second section reports results for the conventional version of the Extended GARCH model with the mean equation given by Equation 4.11 (Extended GARCH model), and the third reports results for the regime dependent version of the final Extended GARCH model (RD-GARCH model). Of

Table 5.5: Descriptive statistics for standardized residuals (final model)

Model	Regime	Mean	Variance	Skewness	Kurtosis	Jarque-Bera	Persistence	Sign Bias	Negative	Positive
model18	none detected	-1.174×10^{-3}	1.000	−0.567	4.913	18522.509	0.96361	2.567	−3.555	−1.403
model22	none detected	-1.595×10^{-4}	1.000	−0.533	4.913	18416.048	0.96247	1.411	−2.941	−0.533
model22	overall	-7.363×10^{-3}	1.000	−0.374	3.259	8145.752	0.95368	0.708	−1.432	−0.013
	base	-4.956×10^{-3}	0.999	−0.513	4.554	8954.209	0.95368			
	low volatility	2.276×10^{-2}	0.997	−0.156	0.804	101.731	0.95811			
	high volatility	-3.556×10^{-2}	1.003	−0.224	2.178	894.999	0.96395			

The table presents additional descriptive statistics for the standardized residuals from the mean equation for the GARCH model selected in Chapter 4. The first row of the table presents the statistics that result when no regime detection and separation is performed. The final four rows present the statistics which result when regimes of unusually high and low volatility are detected and the parameters of the GARCH model are allowed to change in the different regimes. The first row of this section (the second row of the table) presents the statistics for the entire set of standardized residuals. The second row of this section (the third row of the table) presents the statistics for the residuals from the base regime only, and the final two rows present the statistics for the low and high volatility regimes respectively. In all rows of the table, the first through fourth columns present the first four statistical moments of the standardized residuals from the mean equation. The fifth column presents the Jarque and Bera (1987) test for normality. The sixth column presents the first order autocorrelation of the square root of the conditional variance of the standardized residuals, which is a measure of the persistence of the conditional volatility. The seventh through ninth columns present the results of the sign-bias tests suggested by Engle and Ng (1993) and explained in Section 3.3 of the text. The sign-bias test results are not reported for the final three rows of the table as the segmented nature of the regime results makes the application of these tests impossible. The data set comprises excess returns on an U.S. equity index covering the period January 2, 1926 to December 29, 1990, a total of 17,484 observations.

particular interest are the kurtosis, persistence, and Sign Bias tests result columns.

The excess kurtosis for the conventional version of the final model is approximately 6.8, which is almost identical to the value from Model 18, indicating that changing the mean equation has not made a difference in the excessively fat tails exhibited by the residuals. Allowing for regime dependence, however, decreases this value to 4.6 overall, and 0.8, 2.2, and 3.3 in the base, low and high volatility periods respectively. The standard error of the kurtosis is again given by $\sqrt{24/N}$, which gives values of 4.93×10^{-2}, 8.56×10^{-2}, and 7.43×10^{-2} for the base, low, and high volatility regimes. Introduction of regime dependence has made a large difference, but the residuals continue to exhibit statistically significant excess kurtosis.

The persistence column from Table 5.5 shows the first order autocorrelation in h_t, and provides an different representation of the persistence in the conditional volatility from that provided by the HIC. The results presented in the table confirm those presented in the HIC that persistence is not dramatically affected by allowing the model parameters to vary across regimes. While the persistence figure reported in Table 5.5 provides similar evidence regarding persistence of conditional volatility to that presented in Figure 5.3, the HIC is preferable as the shape of the curve can potentially provide additional information regarding the difference in the nature of the persistence of large and small values for conditional volatility. Although it does not appear to be the case for the data used in this section, it is possible that the degree of persistence in conditional volatility might be dependent upon its magnitude; large values of conditional volatility might exhibit a different degree of persistence than small values. The single value for the persistence presented in the table would not reflect this situation.

The final three columns from Table 5.5 present the sign bias tests introduced in Chapter 3. The results here are particularly dramatic. For the Extended GARCH-M model, the sign bias test statistic and the negative sign bias test statistic are both significant, indicating that the sign of the previous shock and the magnitude of the previous negative shock both retain some predictive power for the current squared standardized residual. In the conventional final Extended GARCH model, only the negative sign bias test statistic is significant, indicating that the augmented mean equation has removed the predictive power of the sign of the previous shock, but is unable to account completely for the predictive power of the magnitude of the previous negative shock. Finally, RD-GARCH results in none of the sign bias test statistics being significant, indicating that modeling regime dependence in the conditional volatility by allowing the conditional volatility parameters to change with regimes has removed the predictive power of the magnitude of the previous negative shock.

5.5 Tests of Alternative Distributions

Following the development in Section 4.4, the NICs and HICs are presented in Figures 5.4 and 5.5 for models assuming the GED and Student-t distribution for disturbances respectively. In addition, the model parameters are presented in Table 5.7

Table 5.6: LR tests for alternative distributions

		H_0: Normal Distribution		
Distribution	H_A	Restriction	LR	p-value
GED	k free	$k = 2$	1044.95	< 0.001
Student-t	k free	$\frac{1}{k} = 0$	1190.57	< 0.001

The table presents the results for likelihood ratio tests for alternative distributions. In each case the test compares a model obtained by a restriction on the tail thickness parameter the model which imposes no restriction. The likelihood ratio statistics are all distributed χ^2 with 1 degree of freedom. The corresponding p-values are included in the final column. The data set comprises excess returns on an U.S. equity index covering the period January 2, 1926 to December 29, 1990, a total of 17,484 observations.

and the additional descriptive statistics in Table 5.8.

The column for parameter k in Table 5.7 presents the tail thickness parameter for the alternative distributions. For both distributions, the tail thickness parameter indicates that a leptokurtic distribution is appropriate for the disturbances. Also in both cases, the relationship between the tail thickness parameter in the conventional version of the model and the parameters which result when tail thickness is allowed to vary across volatility regimes is similar—the degree of leptokurtosis is greatest for the base and high volatility regimes, and is somewhat smaller for low volatility regimes.

As the GED and Student-t distribution both nest the normal distribution, a likelihood ratio test can be performed to determine if allowing for either of the alternative distributions has provided an improvement in model fit. Neither the GED or the Student-t distribution nests the other distribution, so LR comparisons between the two alternatives is not possible. In addition, the parameterization which results from the stepwise estimation procedure used in the regime dependent version of the model means that the regime dependent leptokurtic models which are selected cannot be compared directly to the particular regime dependent normal model with the LR test. Tests for model improvement were conducted on the likelihood values from the conventional model presented in Table 5.4 for the normal distribution and Table 5.7 for the leptokurtic distributions, and the results are presented in Table 5.6. The likelihood ratio statistics are distributed χ^2 with one degree of freedom, and clearly reject the null hypothesis of no difference between the model which assumes normally distributed disturbances and either of the models which assume a leptokurtic distribution. This indicates that either the model with disturbances from the GED or the model with Student-t distributed disturbances provides a better fit than the model with normally distributed disturbances.

Estimates for the parameters from the conditional volatility equation provided in Table 5.7 are generally similar to those presented in Table 5.4 for the final Extended GARCH model. While there are differences in some of the parameters across the models estimated under the different distributional assumptions, it is not possible to provide interpretation of the differences in any particular parameter independent of the other parameters which have changed. The NICs, which allow this comparison

Figure 5.4: News Impact Curve (alternative distributions)

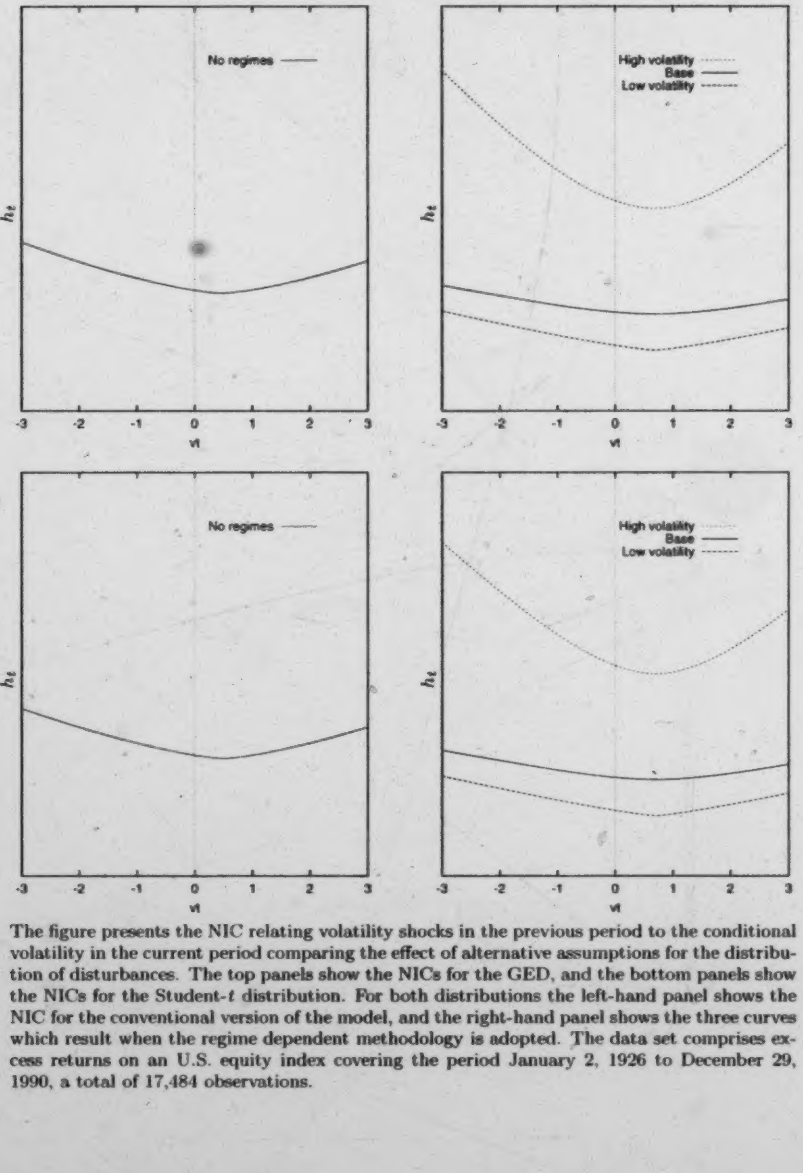

The figure presents the NIC relating volatility shocks in the previous period to the conditional volatility in the current period comparing the effect of alternative assumptions for the distribution of disturbances. The top panels show the NICs for the GED, and the bottom panels show the NICs for the Student-t distribution. For both distributions the left-hand panel shows the NIC for the conventional version of the model, and the right-hand panel shows the three curves which result when the regime dependent methodology is adopted. The data set comprises excess returns on an U.S. equity index covering the period January 2, 1926 to December 29, 1990, a total of 17,484 observations.

across distributional assumptions for all parameters simultaneously, are essentially the same for all three distributional assumptions, and suggest identical interpretation to that offered in Section 5.4. The case for the HICs is similar, although there appears to be slightly less difference between the levels of persistence among the regimes for the leptokurtic distributions than among the regimes for the normal distribution.

Within each of the three distributional assumptions, the result noted in Section 5.4

Figure 5.5: History Impact Curve (alternative distributions)

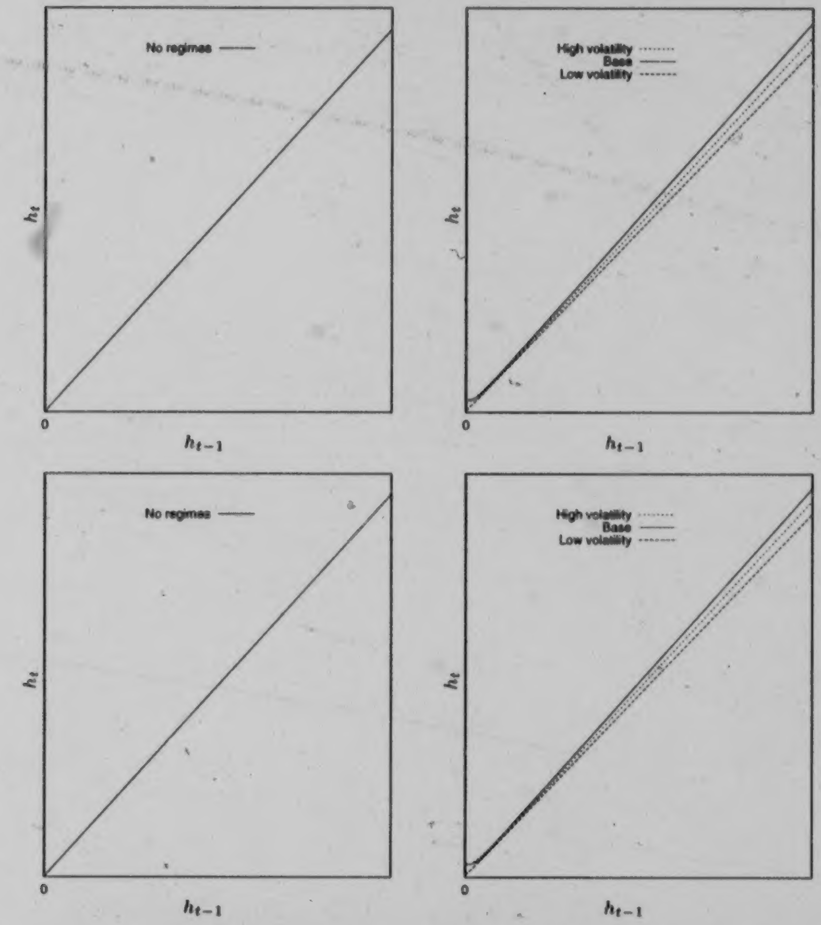

The figure presents the HIC relating the conditional volatility in the previous period to the conditional volatility in the current period comparing the effect of alternative assumptions for the distribution of disturbances. The top panels show the HICs for the GED, and the bottom panels show the HICs for the Student-t distribution. For both distributions the left-hand panel shows the HIC for the conventional version of the model, and the right-hand panel shows the three curves which result when the regime dependent methodology is adopted. The data set comprises excess returns on an U.S. equity index covering the period January 2, 1926 to December 29, 1990, a total of 17,484 observations.

that almost no difference in the mean equation parameter estimates results from allowing the conditional variance to depend upon the regime continues to hold. The parameters from the mean equation change across different distributions, however. In particular, allowing for a leptokurtic distribution of disturbances has caused a slight decrease in the magnitude of the impact of the return from the prior period, and slight increase in the (negative) magnitude of the return from two periods ago.

Table 5.8 presents the additional diagnostic statistics for the models obtained under the alternative assumptions for distributions and may be compared to the results presented in Table 5.5 for the model which assumes normally distributed disturbances. Derivation of the likelihood functions for the alternative distributions are presented in Appendix C. In each case, the distributional assumption is defined in terms of the standardized residual, v_t, with a change of variables transforming the density function to provide the density of r_t, the return series. For this reason the standardized residuals will, by construction, not be normally distributed. As a result, Table 5.8 does not present the Jarque-Bera test statistic for normality, and discussion of the normality of the standardized residuals in not appropriate.

Comparison of the persistence statistics from the tables generally confirms the observations made in reference to the HICs that the degree of persistence in conditional volatility does not change notably under the different distributions, although it does systematically appear to be slightly higher.

The sign bias test statistics from the leptokurtic models are particularly interesting when compared to those from the model which assumes normally distributed disturbances. In all cases, the sign bias statistics for the models with leptokurtic disturbances are larger than the corresponding statistics for the model with normally distributed disturbances. In fact, the negative sign bias test statistic which is not significant in the model which assumes normally distributed disturbances is significant in both the model which assumes disturbances from the GED and the model which assumes disturbances from the Student-t distribution. This indicates that the assumption of normality for disturbances provides greater independence for the standardized residuals than does either of the alternative distributions.

The results from the sign bias tests suggest a much different course of action than that suggested by the results of the likelihood ratio tests. The likelihood ratio tests indicate that either the GED or the Student-t distribution would be appropriate in the case of conventional models, but are unable to provide an indication as to which of the two distribution would be preferred and cannot provide information regarding the regime dependent versions of the models. The sign bias tests indicate that an assumption of normality for disturbances is appropriate in the case of regime dependent models, and do not suggest that one distribution is preferred over any other for the conventional version of the model. In light of this lack of a clearly preferred choice for a distributional assumption regarding the disturbances, the assumption of normality will be retained for the tests of robustness which appear in Chapter 6, and the question of the most appropriate distribution for prediction purposes will be revisited in Chapter 7 where the performance of the various distributional assumptions at out-of-sample forecasting will be addressed.

Table 5.7: Estimates of model parameters (alternative distributions)

Regime	k	γ_0	γ_2	γ_3	γ_4	λ	ν	b	ω	α	β	Likelihood
Generalized Error Distribution												
none detected	1.308 (0.029)*	$7.902^{[-4]}$ $(5.434^{[-5]})$*	0.143 (0.009)*	−0.051 (0.008)*	−0.002 $(1.340^{[-4]})$*	1.020 (0.266)*	1.437 (0.094)*	0.478 (0.038)*	$1.092^{[-4]}$ $(1.412^{[-4]})$	0.069 (0.017)*	0.920 (0.017)*	59615.531
base	1.327 (0.034)*	$7.758^{[-4]}$ $(5.712^{[-5]})$*	0.144 (0.009)*	−0.050 (0.008)*	−0.002 $(1.374^{[-4]})$*	3.792 (10.427)	1.866 (0.928)*	0.686 (0.049)*	$2.962^{[-10]}$ $(1.535^{[-8]})$	0.124 (0.282)	0.806 (0.374)*	59771.255
low volatility	0.188 (0.065)*					−2.705 (7.301)	−0.618 (0.210)*	0	$2.026^{[-4]}$ (0.003)	0	0	
high volatility	0					−1.890 (5.330)	0	0	$1.089^{[-5]}$ $(2.458^{[-4]})$	0	0	
Student-t Distribution												
none detected	6.062 (0.294)*	$7.653^{[-4]}$ $(5.407^{[-5]})$*	0.151 (0.008)*	−0.049 (0.008)*	−0.002 $(1.315^{[-4]})$*	1.097 (0.304)*	1.398 (0.091)*	0.501 (0.032)*	$7.335^{[-5]}$ $(1.090^{[-4]})$	0.077 (0.020)*	0.914 (0.020)*	59688.341
base	6.235 (0.339)*	$7.518^{[-4]}$ $(5.360^{[-5]})$*	0.151 (0.008)*	−0.048 (0.007)*	−0.002 $(1.249^{[-4]})$*	4.144 (0.870)*	1.846 (0.130)*	0.702 (0.036)*	$4.922^{[-11]}$ $(2.135^{[-10]})$	0.145 (0.031)*	0.784 (0.038)*	59829.310
low volatility	3.266 (1.446)*					−2.903 (0.692)*	−0.601 (0.173)*	0	$8.975^{[-5]}$ $(1.825^{[-4]})$	0	0	
high volatility	0					−1.940 (0.488)*	0	0	$2.731^{[-6]}$ $(6.277^{[-6]})$	0	0	

Entries of the form $x^{[-y]}$ represent scientific notation; that is $2.5^{[-2]}$ represents the value 2.5×10^{-2}.

The table presents the parameter estimates for the GARCH model selected in Chapter 4 comparing the effect of alternative distributions of disturbances. Numbers in parentheses under the parameter estimates represent asymptotic standard errors. Parameters which are significantly different from zero at the 5% level are indicated with an asterisk. The table is broken into two sections. The first section presents the results for the generalized error distribution, and the second section presents the results for the Student-t distribution. Within each section, the first two rows of are the parameter estimates and corresponding standard errors that result when no regime detection and separation is performed. The remaining six rows present the parameter estimates and standard errors which result when regimes of unusually high and low volatility are detected and the parameters of the GARCH model are allowed to change in the different regimes. Parameters which are excluded by the stepwise estimation procedure are represented by a value of zero, with no corresponding standard error presented. The data set comprises excess returns on an U.S. equity index covering the period January 2, 1926 to December 29, 1990, a total of 17,484 observations.

Table 5.8: Descriptive statistics for standardized residuals (alternative distributions)

Regime	Mean	Variance	Skewness	Kurtosis	Persistence	Sign Bias	Negative	Positive
GED								
none detected	$-.2807 \times 10^{-2}$	1.015	−0.554	5.129	0.96582	1.581	−3.140	−0.742
overall	-2.591×10^{-2}	1.010	−0.407	3.712	0.97296	1.090	−2.061	−0.260
base	-2.440×10^{-2}	1.008	−0.574	5.338	0.97296			
low volatility	-5.301×10^{-3}	1.001	−0.157	0.820	0.95372			
high volatility	-4.488×10^{-2}	1.021	−0.216	2.212	0.96476			
Student-t Distribution								
none detected	$-.2552 \times 10^{-1}$	1.011	−0.558	5.197	0.96709	1.432	−3.251	−0.603
overall	-2.346×10^{-2}	1.007	−0.413	3.809	0.97282	0.827	−2.108	−0.150
base	-2.190×10^{-2}	1.006	−0.579	5.474	0.97282			
low volatility	-1.318×10^{-3}	0.996	−0.151	0.827	0.95427			
high volatility	-4.370×10^{-2}	1.018	−0.233	2.268	0.96804			

The table presents additional descriptive statistics for the standardized residuals from the mean equation for the GARCH model selected in Chapter 4 comparing the effect of alternative distributions of disturbances. The table is broken into two sections. The first section presents the statistics for the generalized error distribution, and the second section presents the results for the Student-t distribution. Within each section the first row presents the statistics that result when no regime detection and separation is performed. The remaining rows present the statistics which result when regimes of unusually high and low volatility are detected and the parameters of the GARCH model are allowed to change in the different regimes. The second row presents the statistics for the entire set of standardized residuals. The third row presents the statistics for the residuals from the base regime only, and the final two rows present the statistics for the low and high volatility regimes respectively. In all rows of the table, the first through fourth columns present the first four statistical moments of the standardized residuals from the mean equation. The fifth column presents the first order autocorrelation of the square root of the conditional variance of the standardized residuals, which is a measure of the persistence of the conditional volatility. The sixth through eighth columns present the results of the sign-bias tests suggested by Engle and Ng (1993) and explained in Section 3.3 of the text. The sign-bias test results are not reported for the final three rows of each section as the segmented nature of the regime results makes the application of these tests impossible. The data set comprises excess returns on an U.S. equity index covering the period January 2, 1926 to December 29, 1990, a total of 17,484 observations.

Chapter 6

Model Stability

Previous chapters have identified a model, proposed several extensions to that model, and analyzed the application of the extended model to a particular dataset. Having developed the model and analyzed the results from its application, it is appropriate to investigate the stability of the results and to determine the extent to which the estimates produced are robust. This chapter considers three topics related to stability and robustness: raw vs. excess returns, stability of the regime selection process, and parameter stability.

6.1 Return Series

The dataset used for model development and the initial analysis comprised the daily excess return on U.S. equities covering the period January 2, 1926 through December 31, 1990. As outlined in Section 3.1.1, the series used for the raw return was a combination of the daily CRSP value-weighted index of U.S. common stocks for

Figure 6.1: Annualized monthly risk-free rate 1926–1990

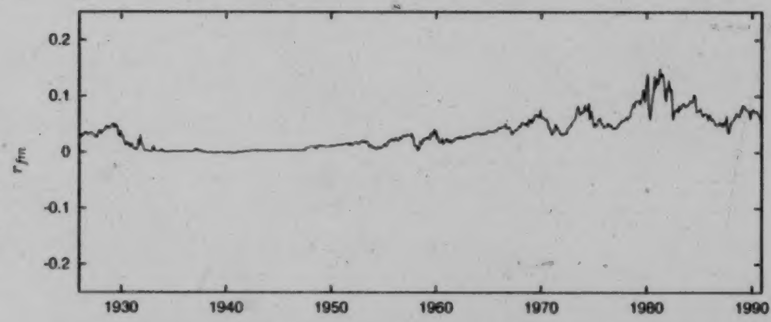

The figure presents the monthly risk-free series covering the period January 2, 1926 through December 31, 1990 (864 observations) where the returns have been converted logs using the formula: $r_{fm} = \ln(1 + R_{fm}/100)$ where R_{fm} is the annualized 1 month average risk-free rate from the Fama/Bliss dataset.

the period from July 3, 1962 through December 31, 1990, and the Schwert (1989) data for the period January 2, 1926 through July 2, 1962. The risk-free rate used to construct the excess return was the Fama/Bliss 1-month average risk-free series from the CRSP bond dataset, which was converted to a daily return series by dividing by the number of calendar days. This particular construction was chosen as it allowed the model development results to use the results reported in Hentschel (1995) as a base reference.

This approach to generating a daily series from monthly data makes the implicit assumption that the risk-free rate is sufficiently stable that a simple straight-line interpolation of the monthly rate provides a reasonable proxy for the daily rate. Figure 6.1 presents the logged monthly risk-free series covering the 1926–1990 period. The returns have been converted to logs as outlined in Section 3.1.1:

$$r_{fm} = \ln(1 + R_{fm}/100) \tag{6.1}$$

where R_{fm} is the annualized one month average risk-free rate from the Fama/Bliss dataset. The figure indicates that the assumption that the risk-free rate is sufficiently smooth to allow straight-line interpolation of daily rates from the monthly rate is likely appropriate for much of the sample period. For the latter part of the period, however, this assumption may be problematic. The risk-free rate exhibits notable volatility following the late 1950's, and in particular is very volatile during the late 1970's and early 1980's. If the risk-free rate in this period is too volatile for the linear interpolation of monthly data to provide a reasonable approximation of the daily rate, construction of the excess return series in this manner will introduce additional noise into the excess return series.

Over the 1926–1990 sample period, the mean magnitude of the daily raw return on the U.S. equity series was 0.71%. Over the same period the mean magnitude of the daily risk-free rate obtained from $r_{fm}/365$ was 0.0090%, nearly two orders of magnitude smaller.[1] This difference in magnitude, coupled with the questionable appropriateness of using the converted monthly risk-free rate during a large portion of the sample period, suggests that the use of the raw return alone might be more appropriate.

6.1.1 Stability of the Regime Selection Process

As development of the model was accomplished with the excess return series, the final model is first re-estimated on the raw return series. A comparison of the regimes detected is presented graphically in Figure 6.2, and in tabular form for the raw return series in Table 6.1[2] and for the excess return series in Table 5.1. Comparison of the tables indicates that the same number of low volatility regimes are detected in both

---

[1] The median magnitude of the raw return was 0.47% and the median magnitude of the risk-free rate was .0073% over this period, again a difference of nearly two orders of magnitude.

[2] Because all tables for this chapter are presented in a sideways orientation, they have been grouped together at the end of the chapter to improve the readability of the text.

Figure 6.2: Returns 1926–1990

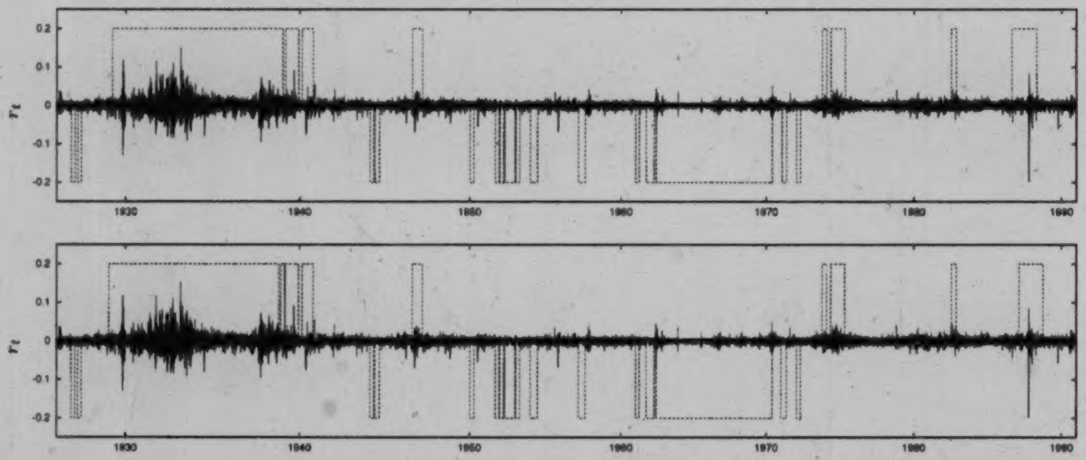

The top panel of the figure repeats Figure 5.1 for comparative purposes, and presents the excess returns on the stock index series for the period January 2, 1926 through December 31, 1990 (17,484 observations). The bottom panel presents the raw returns on the same stock index for the same period. In both panels, regimes of unusually high or low volatility in the returns are indicated by boxes above or below the zero axis respectively.

datasets, with 11 of the 16 regimes sharing identical beginning and ending dates, and all 16 having at least one terminal date in common. A total of 27 of the 32 terminal date are common to both datasets. In four of the five remaining cases, the difference is three or fewer trading days. For the final regime, which occurred in early 1971, the regime detected in the raw return series is 17 trading days longer than the regime detected in the excess return series. The excess return series has 3279 days in low volatility regimes and the raw return series has 3304 days. The high volatility regimes are somewhat less similar. Eight high volatility regimes are detected in the excess returns series, and nine are detected in the raw returns series. The first two high volatility regimes in the raw return series correspond to the first high volatility regime in the excess return series with a period of 28 trading days from October 25 to December 5, 1938 detected as belonging to the base regime in the raw returns series. Six of the remaining seven high volatility regimes are characterizes by beginning and ending dates which are within 5 trading days of each other. The final high volatility regime is quite different between the two datasets, beginning 130 trading days later and ending 114 trading days later in the raw return series.

In all, 24 regimes of high and low volatility have been detected in the excess returns series and 25 have been detected in the raw returns series. Twenty-one of these regimes are either identical or within five trading days of being identical, and one regime differs by 17 trading days. In all 22 cases either the beginning or ending date of the regimes are identical between the series. The remaining two regimes differ between the series. The first regime which differs markedly between the two datasets is coincidental with the crash of 1929 and the subsequent long period of extreme volatility characterizing the 1930's. Given the magnitude of the volatility during this period, the difference between the beginning and ending dates of the regime in the excess return series and the two regimes which cover the same event in the raw return series is not remarkable. The second regime which differs notably between the two series is coincidental with the market crash of 1987. As this regime is much smaller than the 1929–1939 regime, the differences between the regimes detected in the two datasets is more troubling. Inspection of the two return series in Figure 6.2 shows that the crash occurred during a period of moderate volatility, and that the high volatility event surrounding the crash itself is quite short-lived. While the high volatility regime detected and shown in either series appears reasonable, the regime which begins and ends later, as detected in the raw return series, does appear to capture the event more appropriately.

6.1.2 Parameter Stability

Table 6.3 presents the parameter estimates from the excess and raw return series for the case where no regime detection and separation is performed. These results show no difference between the parameter estimates obtained with the excess returns series and those obtained with the raw returns series. While numerical differences do exist between the two sets of parameter estimates, the differences are small and there are no changes in the significance of the parameter estimates.

Table 6.4 presents the results for the case where regimes are detected and the

Figure 6.3: News and History Impact Curves 1926–1990 (excess returns)

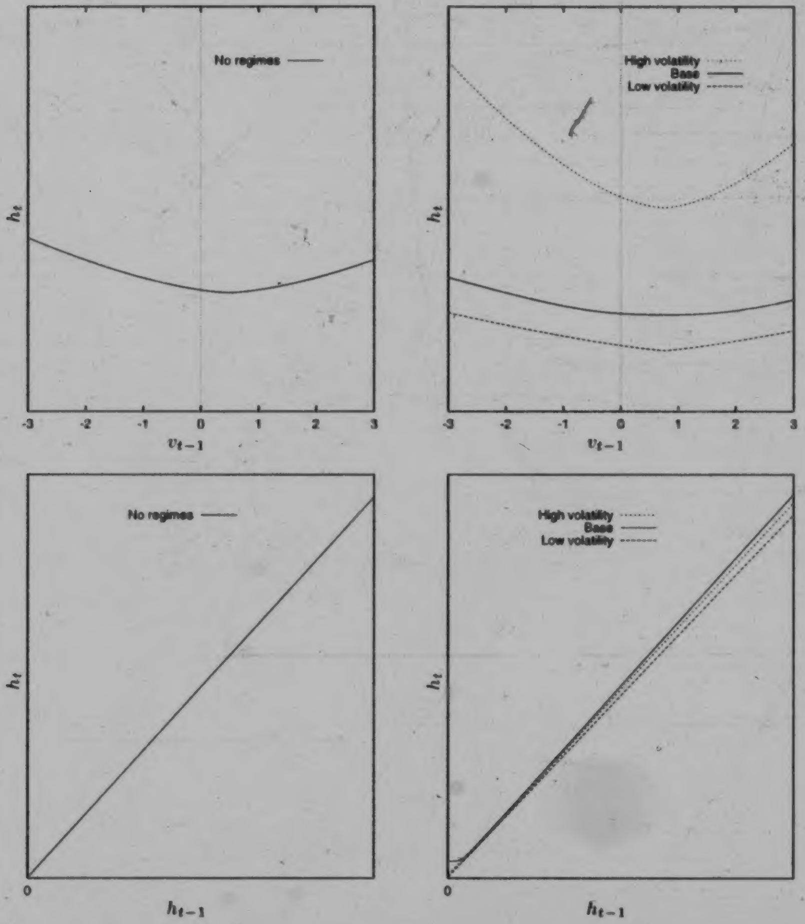

The top two panels of the figure present the NIC relating volatility shocks in the previous period to the conditional volatility in the current period for the excess return series covering the dates January 1, 1926 through December 31, 1990 (17,484 observations). The left-hand panel shows the NIC for the conventional version of the model, and the right-hand panel shows the three curves which result when the regime dependent methodology is adopted. The range of volatility shocks over which the figure is presented corresponds to the range ±3. Since $v_t \sim iid(0,1)$ this corresponds to ±3 standard deviations. Both figures include a dashed vertical line corresponding to $v_{t-1} = 0$ for comparison of asymmetry effects. The bottom two panels of the figure present the HIC relating the conditional volatility in the previous period to the conditional volatility in the current period for the excess return series covering the same period. The left-hand panel shows the HIC for the conventional version of the model, and the right-hand panel shows the three curves which result when the regime dependent methodology is adopted. To facilitate comparison, the vertical scale of each pair of figures is the same.

Figure 6.4: News and History Impact Curves 1926–1990 (raw returns)

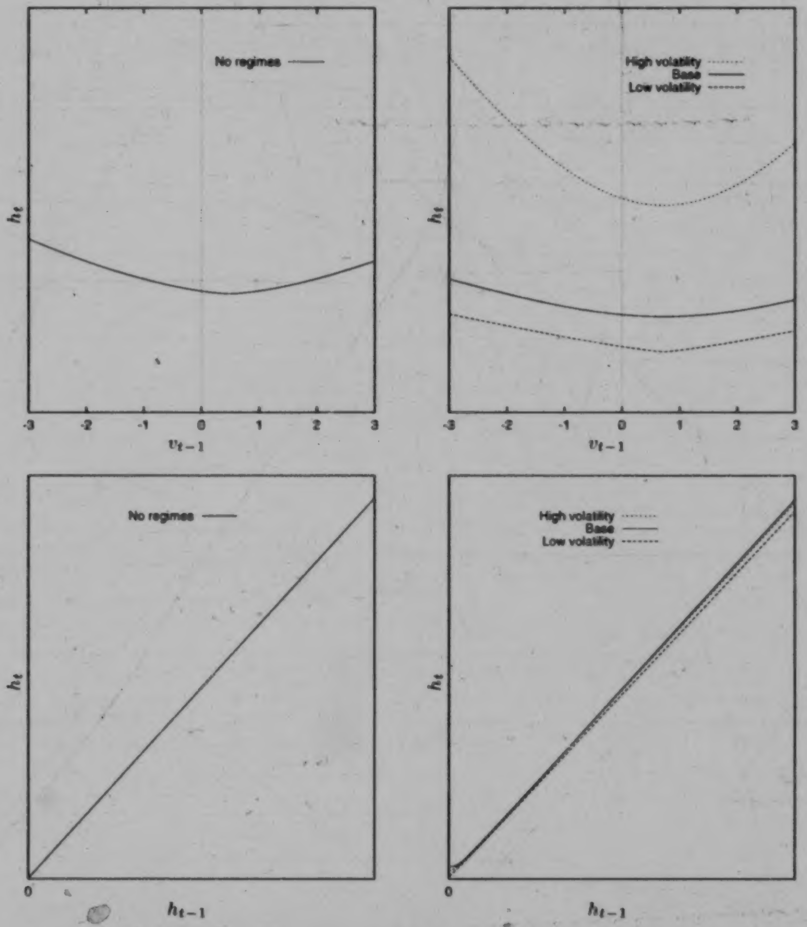

The top two panels of the figure present the NIC relating volatility shocks in the previous period to the conditional volatility in the current period for the raw return series covering the dates January 1, 1926 through December 31, 1990 (17,484 observations). The left-hand panel shows the NIC for the conventional version of the model, and the right-hand panel shows the three curves which result when the regime dependent methodology is adopted. The range of volatility shocks over which the figure is presented corresponds to the range ±3. Since $v_t \sim iid(0, 1)$ this corresponds to ±3 standard deviations. Both figures include a dashed vertical line corresponding to $v_{t-1} = 0$ for comparison of asymmetry effects. The bottom two panels of the figure present the HIC relating the conditional volatility in the previous period to the conditional volatility in the current period for the raw return series covering the same period. The left-hand panel shows the HIC for the conventional version of the model, and the right-hand panel shows the three curves which result when the regime dependent methodology is adopted. To facilitate comparison, the vertical scale of each pair of figures is the same.

parameter values are permitted to vary across regimes. There are some differences between the parameter estimates obtained with the excess return series and those obtained with the raw return series when regime dependence is permitted. The parameters which have been excluded by the stepwise estimation procedure also differ. It is difficult, however, to asses the impact of the differences between the parameter estimates since an increase in λ can, for example, be partially offset by decreases in either ω or α. For this reason, the NICs and HICs are presented in Figure 6.3 for the excess return series and in Figure 6.4 for the raw return series. Comparison of the NICs presented in these sets of figures indicate little detectable difference in the impact of the prior period shock on current conditional volatility between the two data series; the various changes in the parameter estimates tend to compensate for each other. There is a small difference in the curvature of the NIC for the high volatility regimes, but this difference is unremarkable. The HICs tell a slightly different story, however. There is less difference in the persistence of the conditional volatility among the regime types for the raw return series than there is for the excess return series. This difference appears to be largely due to a decrease in the persistence of the conditional volatility in the base regime for the raw returns series.

As there does not appear to be a notable difference in the results obtained with the raw return series and those obtained with the excess return series, and as the raw return series avoids the problems associated with using a monthly risk-free series to construct a daily excess return series, the subsample comparisons in the next section will be conducted on the raw return series.

6.2 Sub-sample Results

An advantages of working with raw returns rather than excess returns is that the data for the raw daily returns from Schwert (1989) is available from February 17, 1885 while the Fama/Bliss monthly risk-free rate is only available from January, 1926. By working with the raw return series, it is possible to test the stability of the model parameters over a larger time horizon. For the purposes of comparing sample periods, several additional time periods will be considered:

The entire sample period is extended to comprise the range January 2, 1915 through December 29, 1995. Although the Schwert data begin in 1885, it is not possible to make use of the entire range of data provided. From July 31, 1914 through December 11, 1914, the New York Stock Exchange was closed due to the effects of the First World War, resulting in a break of $4\frac{1}{2}$ months in the dataset. As a result, the first trading day of 1915 has been selected as the beginning of the enlarged sample period. The CRSP data which comprise the equity index series after July 3, 1962 is currently available to December 31, 1997. The tests of stability and robustness are conducted with a sample period ending December 29, 1995. This provides a two year hold-out sample for testing the out-of-sample properties of the RD-GARCH methodology.

The 1915–1995 sample period is then broken into two sub-periods: January 2, 1915 through December 31, 1951 and January 2, 1952 through December 29, 1995.

The choice of the end of 1951 as the break point for the subsamples is made following Campbell and Hentschel (1992), and serves to separate the period of extremely high volatility which corresponds to the Great Depression from the postwar period.

Finally, the 1915–1951 and 1952–1995 periods are each further divided into two sub-periods. The 1915–1951 period is broken at the end of the third quarter of 1929, which further isolates the period of extreme volatility associated with the Great Depression from the period surrounding the First World War, and represents a point in time at which the economy of the United States begins to become a more dominant force internationally. The 1952–1995 period is broken at the end of 1973, which, with the oil crisis, represents another shift in the international positioning of the economy of the United States. These additional sub-periods allow additional assessment of the stability of the model through time, and will also allow for the recent behaviour of the market to be assessed more fully. The most recent sub-periods will be used in Chapter 7 to assess the impact of different sample period lengths on the performance of out-of-sample forecasts.

6.2.1 Stability of the Regime Selection Process

Table 6.1 presents the dates for regimes of unusually high and low volatility detected in the raw returns for the 1926–1990 sample period used in the development and initial testing of the RD-GARCH model. Table 6.5 presents the same information for the enlarged sample period beginning January 2, 1915 and ending December 29, 1995. A total of 16 low volatility regimes and nine high volatility regimes were detected in the 1926–1990 sample period, compared to 14 low volatility and 10 high volatility regimes in the portion of the 1915–1995 period which overlaps the smaller sample.

Of the low volatility regimes, 10 regimes are either identical or begin and end within a week of each other in both samples. There are three regimes which have either a beginning date or an ending date, but not both, which is different by more than a week in the two periods, and a single regime for which both terminal dates are more than a week apart in the two samples. In this final case, the ending dates are only eleven days apart. Finally, there are two regimes covering the periods January 6, 1944 through March 24, 1944 (78 days) and January 5, 1961 through April 4, 1961 (89 days) which are detected in the shorter sample but not in the longer sample. In all cases, the regimes detected in the shorter (1926–1990) sample are either the same length as or longer than those detected in the longer (1915–1995) sample.

The results for the comparison of the high volatility regimes are similar. Of the high volatility regimes, four regimes are either identical or begin and end within a week of each other in both samples. There are a total of two regimes detected in the longer sample and three in the shorter sample which cover the period of extreme volatility of the 1930's. The two regimes in the longer sample span the period November 19, 1928 through December 11, 1939 (4039 days) with a nine day period from May 9 to 18, 1939 which is detected as belonging to the base regime. The three regimes in the shorter sample span the period January 9, 1929 through December 11, 1939 (3988 days) with two short periods, October 25 to December 5, 1938 (41 days) and February 28 to March 9, 1939 (nine days), which are detected as belonging to the base regime. Overall these two groups of regimes have a beginning date which is separated by 51 days, and share a common ending date. One of the remaining two regimes which are detected in both sample periods shares a common ending date and has beginning dates which are separated by 12 days. The second regime has ending dates separated by 21 days and beginning dates separated by 62 days. Finally, there are two high volatility regimes covering the periods September 1 to November 30, 1955 (90 days) and March 21 to June 18, 1962 (89 days) which are detected in the longer sample but not in the shorter sample. In all cases, the regimes detected in the shorter (1926–1990) sample are shorter than those detected in the longer (1915–1995) sample.

The similarities and differences between the regimes detected in these two sample periods suggest two summary observations. Given the fact that the longer sample covers 16 additional years or 4554 additional observations, the level of agreement between the regimes detected in the overlap period is remarkable. The nature of the differences between the regimes detected in the two samples is also interesting. The shorter sample provided longer low volatility regimes, and also provided two additional low volatility regimes. The longer sample provided longer high volatility regimes, and also provided two additional high volatility regimes.

A possible explanation for the differences between the regimes detected in the two samples is related to the sequential nature of the regime detection and isolation procedure outlined in Section 4.2.3 and Appendix B. The base regime is defined to be any portion of the sample which has not yet been determined to be part of a regime of unusual volatility. The period of extremely high volatility which occurred during the period following the Crash of 1929 in particular tends to increase the volatility of the base regime substantially prior to its removal by the regime detection algorithm. As this period of high volatility represents a larger portion of the total sample for the shorter sample period than it does for the longer period, the effect will be more pronounced for the shorter sample. In essence, the presence of the extremely high volatility observed in the 1930's raises the threshold for the detection of low volatility regimes, causing a relative bias in favour of longer regimes of low volatility or additional low volatility regimes in the shorter sample period. A similar argument suggests the presence of longer or additional high volatility regimes in the longer sample.

A final note concerns the additional high volatility regime covering the period March 21 to June 18, 1962 in the longer sample period. In both sample periods, the first regime detected and removed is the low volatility period which occurred during the 1960's. For the shorter sample period, this regime encompasses the dates May 29, 1962 to June 1, 1970 while for the longer sample period the corresponding date range is June 19, 1962 to May 28, 1970. The higher volatility corresponding to the March–June 1962 period appears to have triggered the end of the low volatility regime in the shorter sample, while being detected as a high volatility regime in the longer sample.

A comparison of the regimes detected in the full 1915–1995 sample with those detected in the 1915–1951 and 1952–1995 subsample periods is presented graphically in Figure 6.5 and in tabular form in Tables 6.5, 6.6, and 6.7. Comparing first the earlier subsample with the same period from the full sample reveals some strong

Figure 6.5: Raw returns 1915–1995

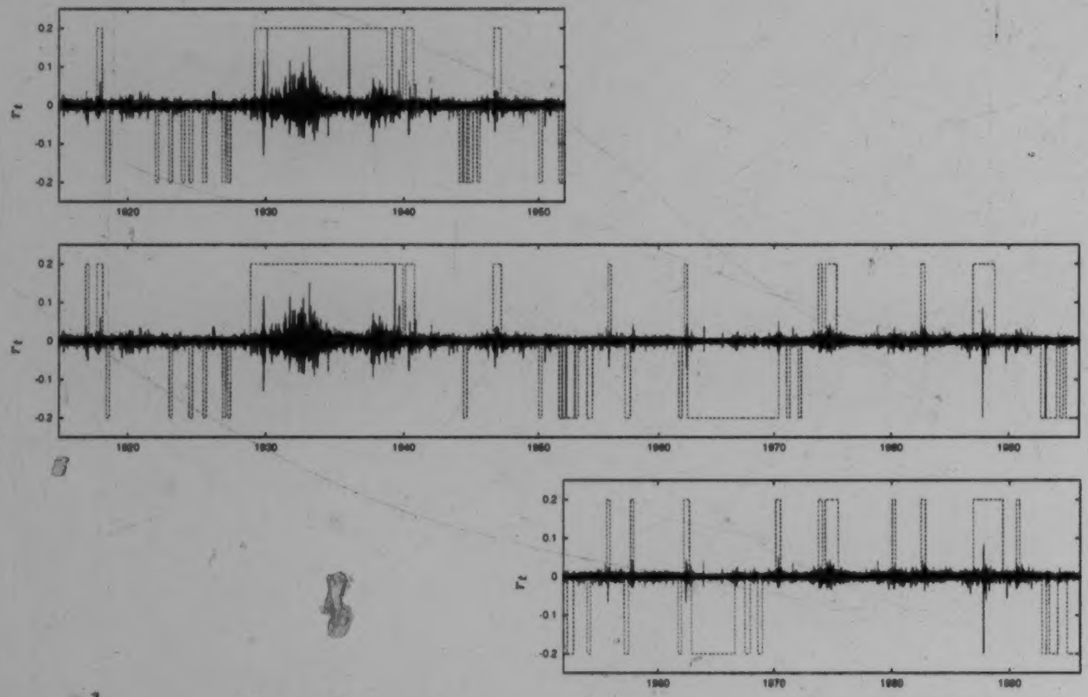

The figure presents the raw returns on the U.S. equity index for the period January 2, 1915 through December 29, 1995 (22,038 observations). The center panel presents the raw returns and regimes detected in the returns for the entire sample period. The top panel presents the returns and regimes detected for the first subsample period which comprises the dates January 2, 1915 through December 31, 1951 (10,944 observations), and the bottom panel presents the returns and regimes detected for the second subsample period which comprises the dates January 2, 1952 through December 29, 1995 (11,092 observations). In all three panels regimes of unusually high or low volatility in the returns are indicated by boxes above or below the zero axis respectively.

similarities as well as some notable differences. The earlier subsample presents 14 low volatility regimes and seven high volatility regimes, compared to nine low and six high volatility regimes in the same period from the full sample. All nine low volatility regimes in the longer sample have corresponding regimes in the shorter sample. In five cases, both the beginning and ending dates are within a week of each other in the two sample periods. A further three regimes have one of the terminal dates within three days of each other and the other within 14 days, and the remaining regime has a beginning date which is different by 8 days and an ending date which differs by 14 days across the two sample periods. The shorter sample period also has five additional low volatility regimes. There is somewhat less correspondence between the sample periods for high volatility regimes. The two regimes associated with the extreme volatility period corresponding to the Great Depression in the larger sample are represented by a total of four regimes in the shorter sample. As outlined above, the two regimes in the longer sample span the period November 19, 1928 through December 11, 1939 (4039 days) with a nine day period from May 9 to 18, 1939 which is detected as belonging to the base regime. The four regimes in the shorter sample span the period March 13, 1929 through December 4, 1939 (3918 days) with three periods, February 15 to 18, 1930 (three days), January 14 to February 10, 1936 (37 days), and October 25, 1938 to March 2, 1939 (128 days) which are detected as belonging to the base regime. Overall these two groups of regimes have a beginning date which is separated by 114 days, and an ending date which is separated by seven days. While none of the remaining regimes have both terminal dates within a week of each other in the two samples, two of the three do have one terminal date within two days of each other, and the other terminal date separated by 15 and 35 days respectively. The final common regime has a beginning date which is different by 16 days and an ending date which is different by 21 days. There is a single high volatility regime which was detected in the longer sample with no corresponding regime in the shorter sample. Overall, the pattern of differences between regimes detected in the two sample periods is similar to that observed in the comparison between the 1915–1995 and 1926–1990 sample periods. In all cases, the shorter sample provided longer or additional low volatility regimes, and the longer sample provided longer and additional high volatility regimes.

Comparison of the later subsample with the full sample reveals less commonality. The later subsample presents 10 low volatility regimes and 10 high volatility regimes, compared to 12 low and six high volatility regimes in the same period from the full sample. Seven of the low volatility regimes in the longer sample have corresponding regimes in the shorter sample. In two cases, both the beginning and ending dates are within a week of each other in the two sample periods. A further four regimes have one terminal date within four days of each other, although the other terminal ranges from a low of 10 to a high of 72 days apart. The final common regime corresponds to the low volatility period of the 1960's, spanning the period June 19, 1962 to May 28, 1970 (2900 days) in the longer sample. In the shorter subsample, this period is detected as three separate regimes, in total spanning the period November 13, 1962 to January 7, 1969 (2247 days) with two rather large periods, July 26, 1966 to June 7, 1967 (316 days) and November 21, 1967 to July 23, 1968 (245 days) which are detected as belonging to the base regime. The longer sample period has four additional low volatility regimes. All six of the high volatility regimes detected in the full sample have corresponding regimes detected in the subsample. Of these six regimes, one has both beginning and ending dates within a week of each other in both sample periods, and an additional four have at least one terminal date within a week of each other. The second terminal date ranges from a low of nine to a high of 86 days apart. The final regime encompasses the crash of 1987, and has a beginning date which differs by 19 days and an ending date which differs by 226 days between the two sample periods. Overall, the pattern of differences between regimes detected in the two sample periods is exactly opposite to that observed in the comparison between the full sample and the earlier subsample. In all cases, the shorter sample provided longer or additional high volatility regimes, and the longer sample provided longer or additional low volatility regimes. The difference appears to be due to two factors: the distortion caused by the extreme volatility period of the 1930's is not present in the shorter subsample, however the long period of low volatility observed in the 1960's becomes more significant. In exactly the same way as the high volatility period of the 1930's causes a relative bias in favour of longer regimes of low volatility or additional low volatility regimes in the shorter sample period, the low volatility period of the 1960's causes a relative bias in favour of longer regimes of high volatility or additional high volatility regimes in the shorter sample period. Again, a similar argument suggests the presence of longer or additional low volatility regimes in the longer sample.

The differences between the tabulated regime detection results for the full sample and the subsamples would seem to indicate that the regime detection process is rather unstable, and generates different regimes depending upon the extent of the sample period. Examination of the graphical representation of the regimes presented in Figure 6.5, however, indicates that while the regimes detected in the earlier subsample do in fact mimic the regimes detected in the full sample remarkably well, those detected in the later subsample are also quite similar those detected in the full sample. Generally the regimes which are detected in the full sample and those which are detected in the later subsample appear to encompass the same events, with the differences being related to the extent of the regime rather than its existence. There is a single high volatility regime which was detected in the full sample which does not appear to correspond to any regime detected in the earlier subsample (the earliest high volatility regime shown in the center panel), and four high volatility regimes which have been detected in the later subsample which do not appear to correspond to any regime detected in the full sample (the second, fourth, seventh, and final high volatility regimes shown in the lower panel). The remaining differences appear to be more related to the length of regimes which have been detected in both samples rather than different regimes entirely.

If the cause of the differences between the regimes detected in the full sample and those detected in the later subsample relates to the period of extreme volatility characterizing the Great Depression, using the regimes detected in the later subsample is likely more appropriate both for model building and for forecasting. Breaking the 1915–1951 and 1952–1995 periods each into two additional sub-periods makes it possible for the later sub-period to determine the stability of the regimes detected in

Figure 6.6: Raw returns 1915–1951

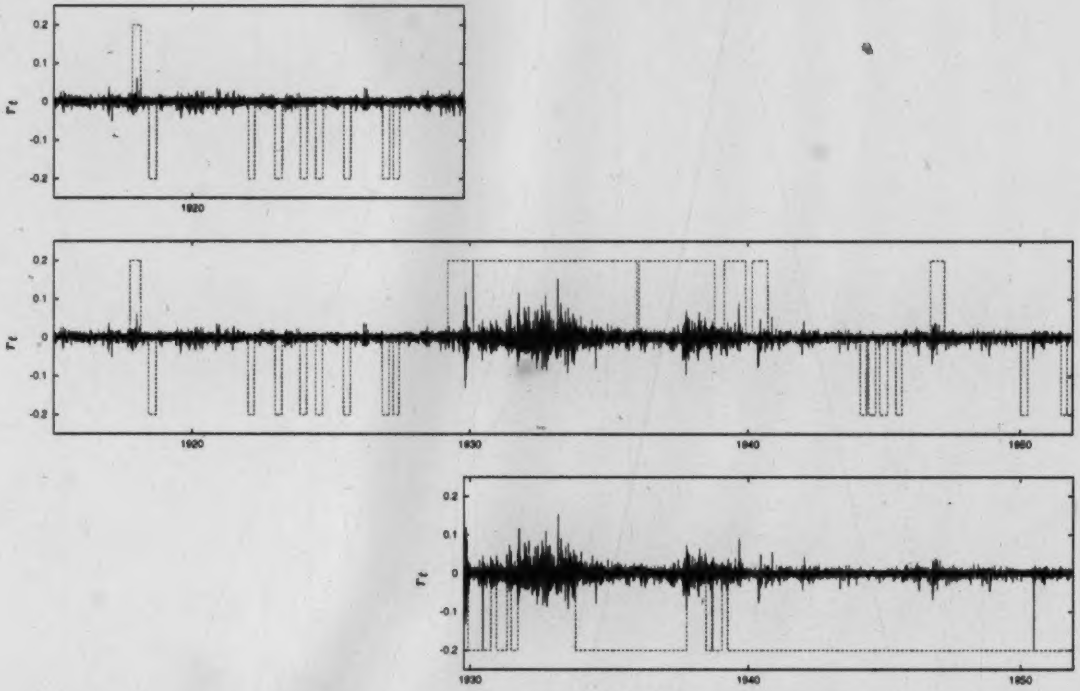

The figure presents the raw returns on the U.S. equity index for the period January 2, 1915 through December 31, 1951 (10,944 observations). The center panel presents the raw returns and regimes detected in the returns for the entire sample period. The top panel presents the returns and regimes detected for the first subsample period which comprises the dates January 2, 1915 through September 30, 1929 (4,406 observations), and the bottom panel presents the returns and regimes detected for the second subsample period which comprises the dates October 1, 1929 through December 31, 1951 (6,536 observations). In all three panels regimes of unusually high or low volatility in the returns are indicated by boxes above or below the zero axis respectively.

Figure 6.7: Raw returns 1952–1995

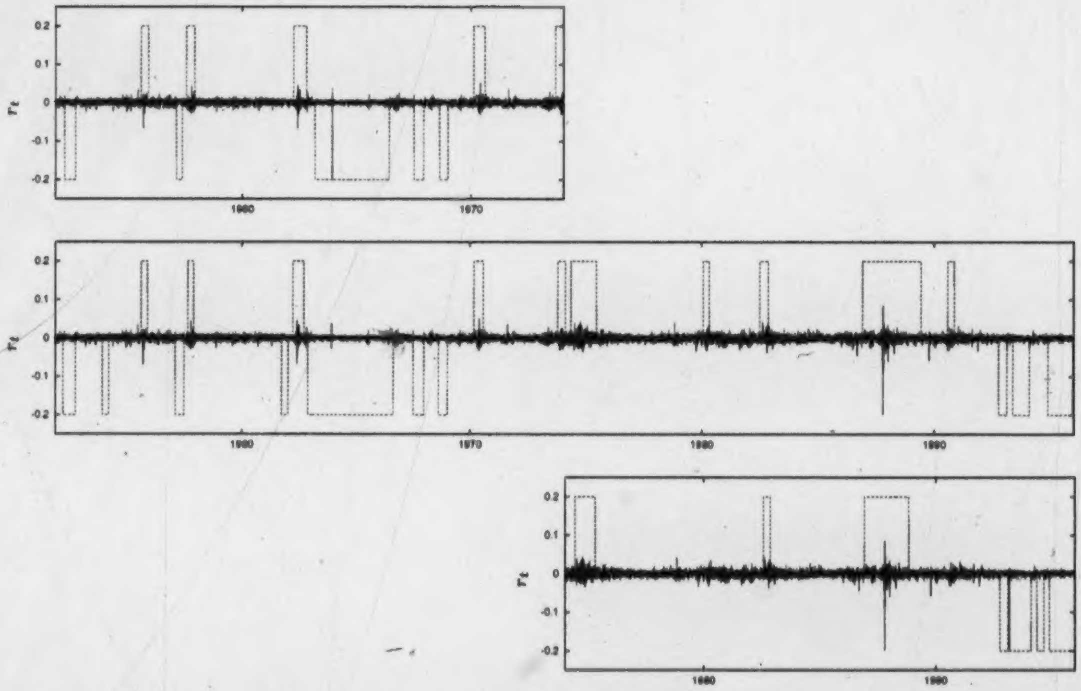

The figure presents the raw returns on the U.S. equity index for the period January 2, 1952 through December 29, 1995 (11,092 observations). The center panel presents the raw returns and regimes detected in the returns for the entire sample period. The top panel presents the returns and regimes detected for the first subsample period which comprises the dates January 2, 1952 through December 31, 1973 (5,531 observations), and the bottom panel presents the returns and regimes detected for the second subsample period which comprises the dates January 2, 1974 through December 29, 1995 (5,559 observations). In all three panels regimes of unusually high or low volatility in the returns are indicated by boxes above or below the zero axis respectively.

the absence of the potential distortions resulting from the inclusion of the extreme volatility period of the 1930's. The results for the application of the regime detection process to the additional four sub-periods are presented in Figures 6.6 and 6.7, and in tabular form in Tables 6.8, 6.9, 6.10, and 6.11.

Comparison of the regimes detected in the 1915–1929 sub-period to those detected in the corresponding date range of the 1915–1951 sub-period shows remarkable agreement. There are a total of eight low volatility regimes detected in the overlapping portion of the two sample periods, of which seven have both beginning and ending dates within a week of each other. The eighth low volatility regime shares a common ending date, and has a beginning date which is different by 10 days. There is a single high volatility regime detected in the overlapping portion of the two sample periods, and it has a beginning and ending dates which are separated by 28 and two days respectively. A likely cause for this degree of agreement between the regimes detected in the two subsamples for the overlapping period is that the extreme volatility period of the 1930's is removed from the longer sample early in the regime detection process. Once this period of extreme volatility is removed, the 1915–1929 period comprises the majority of the remaining data, and the regime detection in the shorter subsample mirrors the detection in the shortened larger sample. The shorter sample provided longer low volatility regimes, and the longer sample provided a longer high volatility regime, mimicking the results observed in the comparison of the full 1915–1995 sample to the 1915–1951 subsample.

As remarkable as the similarity in the results for the 1915–1929 sub-period and the corresponding portion of the 1915–1951 sub-period is the difference between the results for the 1929–1951 sub-period and the latter portion of the 1915–1951 sub-period. In the smaller sample period, the extreme volatility period of the 1930's is detected as belonging to the base volatility regime, and the remaining portion of the sample is detected as being a low volatility regime. There are no high volatility regimes detected in the short subsample. While comparison between the two graphical representations presented in Figure 6.6 can be made if the low volatility regime in the 1929–1951 subsample is compared to the base regime in the larger subsample, there is no way to compare periods of lower than usual volatility. This problem is further demonstrated in Table 6.9 where the variance of the low volatility regimes demonstrates considerably more variability than does the variance of either the high or low volatility regimes in any of the other sub-periods. Taken together, these observations would indicate that the regime detection methodology is likely not appropriate in the case where a relatively short sample period is dominated by a period of extreme volatility.

The results obtained by breaking the 1952–1995 period into two additional sub-periods are not nearly as remarkable as those observed for breaking the 1915–1951 sub-period. The tabulated results do continue to show differences between the regimes detected in the larger 1952–1995 sample and those detected in the 1952–1973 and 1974–1995 subsamples. In general, the 1952–1973 sub-period is characterized by shorter and fewer low volatility regimes and longer high volatility regimes than the corresponding portion of the 1952–1995 subsample. In contrast, the 1974–1995 sub-period is characterized by longer and additional low volatility regimes, and shorter and fewer high volatility regimes than the corresponding portion of the 1952–1995 subsample.

Because the 1952–1973 and 1974–1995 sub-periods are much shorter, the removal of the observations associated with regimes of unusual volatility has a larger impact on the volatility of the base regime. This factor is likely responsible for the differences between the regimes detected in the 1952–1995 subsample and those detected in the 1952–1973 and 1974–1995 subsamples appearing to be more remarkable than the differences between the regimes detected in the full 1915–1995 sample and those detected in the 1915–1951 and 1952–1995 subsamples. Examination of the graphical representation of the detected regimes presented in Figure 6.7 reinforces this observation as the dates and extents of the detected regimes are in fact similar. The systematic differences noted between the two smaller subsamples and the 1952–1995 subsample are consistent with the observation that the overall volatility during the 1952–1973 subsample is somewhat smaller than that during the 1974–1995 subsample.

As was the case with the comparison of the full 1915–1995 sample with the 1915–1951 and 1952–1995 subsamples, the regimes which are detected in the 1952–1995 sample and those which are detected in the 1952–1973 and 1974–1995 subsamples appear to encompass the same events, with the differences being related to the extent rather than the existence of the regimes. There are a total of four regimes detected in the larger sample which do not have corresponding regimes detected in the smaller samples. The second and fourth low volatility regimes shown in the center panel of Figure 6.7 represent low volatility regimes which have been detected in the long sample, but which have no corresponding regimes in the shorter subsample shown in the top panel of the figure. Similarly, the seventh and last high volatility regimes shown in the center panel reflect high volatility regimes which have been detected in the long sample, but which have no corresponding regimes in the shorter subsample shown in the bottom panel of the figure. The remaining differences appear to be more related to the length of regimes which have been detected in both the shorter sample and the corresponding portion of the longer sample rather than different regimes entirely.

In general, with the exception of the 1929–1951 sub-period dominated by the extreme volatility period of the 1930's, the regime detection procedure adopted for this work appears to be quite stable and locates essentially the same regimes regardless of the extent of the sample period selected. As the volatility observed during the Great Depression of the 1930's is unique, and its occurrence is sufficiently far in the past that the domination of the regime selection process is likely not to be an issue for any forecasting applications, the sensitivity of the process to the high volatility regime is likely not material.

6.2.2 Parameter Stability

Conventional Models

Table 6.12 presents the parameter estimates which result when the version of the model which does not allow regime detection and separation is applied to the raw return series. The first section of the table presents the results for the entire extended sample which encompasses the dates January 2, 1915 through December 29, 1995. The second section of the table presents the results for the two subsamples which cover the date ranges January 2, 1915 through December 31, 1951 and January 2, 1952 through December 29, 1995. The final section of the table presents the results for the four subsamples covering the date ranges January 2, 1915 through September 30, 1929, October 1, 1929 through December 31, 1951, January 2, 1952 through December 31, 1973, and January 2, 1974 through December 29, 1995.

Considering first the overall sample covering the period 1915–1995, and the two subsamples created by dividing this period at 1952, several observations obtain: Two of three parameters from the conditional mean equation, γ_2 and γ_3 exhibit notable differences across the sample periods. The effect of the return from the prior period, γ_2, is positive and significantly different from zero in all cases. It takes on a value of 0.145 for the full sample, a value of 0.078 in the earlier subsample, and 0.214 in the later subsample. The parameter value observed for the 1926–1990 subsample, presented in the second section of Table 6.3, is 0.165. The effect of the return two periods ago, γ_3, exhibits similar behaviour. It is negative and significantly different from zero in both the full sample and the later subsample, taking on values of -0.023 and -0.039 respectively. It is not significantly different from zero in the 1915–1951 subsample. The parameter value observed for the 1926–1990 subsample was -0.036. The coefficient of the Monday dummy, γ_4, is significantly different from zero with a value of -0.002 in all three sample periods, as well as in the 1926–1990 subsample.

These changes in autocorrelation structure are further illustrated in the four additional sub-periods presented in the final section of Table 6.12. The lagged return, γ_2, is significantly different from zero and takes on values of 0.088, 0.087, 0.273, and 0.161 in the sub-periods 1915–1929, 1929–1951, 1952–1973, and 1974–1995 respectively. The return lagged two periods, γ_3, exhibits somewhat similar behaviour, becoming significantly different from zero in the 1929–1951 sub-period with a value of -0.033, increasing in magnitude to a value of -0.073 in the 1952–1973 sub-period, and becoming not significantly different from zero in the final 1974–1995 sub-period. In addition, the Monday dummy, γ_4, is significantly different from zero in the three sub-periods prior to 1974, increasing in magnitude from a value of -0.001 in the 1915–1929 sub-period to a value of -0.002 in both the 1929–1951 and 1952–1973 sup-periods. It is not significantly different from zero in the final 1974–1995 sub-period.

The source of the apparent changes in the autocorrelation structure may in fact be due to changes in the construction of the data series itself rather than changes in the structure of the market. As outlined in Section 3.1.1, the data upon which this work is based come from two sources: for the period prior to July 2, 1962, returns are from Schwert (1989); for the period since July 3, 1962, returns are from the CRSP value weighted stock indices for NYSE and AMEX combined. Prior to 1928, the Schwert index is based on a weighted average of the Dow Jones Industrial and Railroad averages which comprised 22 stocks in the beginning, rising through time to 50 stocks. Beginning in 1928 the index is based on the S&P 90 index, changing to the S&P 500 index in 1957. The CRSP index, which replaces the Schwert index in 1962, is a value-weighted index of all issues on the NYSE and AMEX exchanges.

Lo and MacKinlay (1990) suggest that nonsynchronous trading of the stocks in a portfolio—both in the sense that the closing price for some stocks in the portfolio will be from considerably earlier in the day than the closing price for other stocks, and in the sense that some stocks may not trade at all on a particular day—contributes to the autocorrelation commonly noted in portfolio returns. If nonsynchronous trading is contributing to the autocorrelation noted in the subsample periods outlined above, an increase in autocorrelation is to be expected following the change from the S&P 90 index to the S&P 500 index in 1957 as some portion of the 410 added stocks are likely to be more thinly traded than the 90 original stocks. A second increase would be expected following the change from the S&P 500 index to the index comprising the full set of stocks trading on the NYSE and AMEX exchanges. The decrease in autocorrelation between the 1952–1973 subsample and the 1974–1995 subsample is then likely attributable to a general increase in trading volume through time, with a corresponding decrease in thinness of trading and hence decrease in the impact of nonsynchronous trading.

The remaining columns of Table 6.12 present the conditional volatility equation parameters for the various subsamples. As there are considerable differences between the parameters estimates which result from the various sub-periods, and as the interpretation of the various parameters in isolation is problematic for reasons outlined earlier, the comparison of the conditional volatility parameters for the subsamples will be conducted based on the NICs and HICs. In each case, the top left panel of the relevant figure presents the NIC for the conventional version of the model, and the lower left panel presents the corresponding HIC.

Comparing first the 1926–1990 sub-period presented in Figure 6.4 with the full 1915–1995 sample presented in Figure 6.8 shows that the conventional version of the model provides NICs which are nearly identical. The two curves exhibit identical curvature, and are only displaced in space slightly, with the curve for the 1926–1990 sub-period being slightly higher than the curve for the 1915–1995 period. This indicates that the unconditional volatility is slightly higher in the smaller sub-sample, but that the nature of the response of conditional volatility to news is the same for both periods. As the extreme volatility period of the 1930's comprises a larger portion of the smaller sample, the observed greater unconditional volatility is expected. The HICs for the two periods are essentially identical, indicating that there is no change in the level or nature of the persistence of conditional volatility.

Comparison of the full 1915–1995 sample period presented in Figure 6.8 with the 1915–1951 and 1952–1995 sub-periods presented in Figures 6.9 and 6.10 respectively, reveals few notable differences. The location in space of the NIC indicates the unconditional volatility of the relevant sub-period. The 1915–1951 sub-period is located somewhat higher than the full 1915–1995 period, and the 1952–1995 sub-period is located somewhat lower, illustrating the relative effects of the period of extreme volatility associated with the Great Depression and the period of low volatility associated with the 1960's. The earlier sub-period exhibits somewhat more curvature than the full sample period and the later sub-period exhibits somewhat less curvature. There is also a slight additional shift in the curve to the right for the earlier sub-period and to the left for the later sub-period. This observation is confirmed by the tabulated

Figure 6.8: News and History Impact Curves 1915–1995

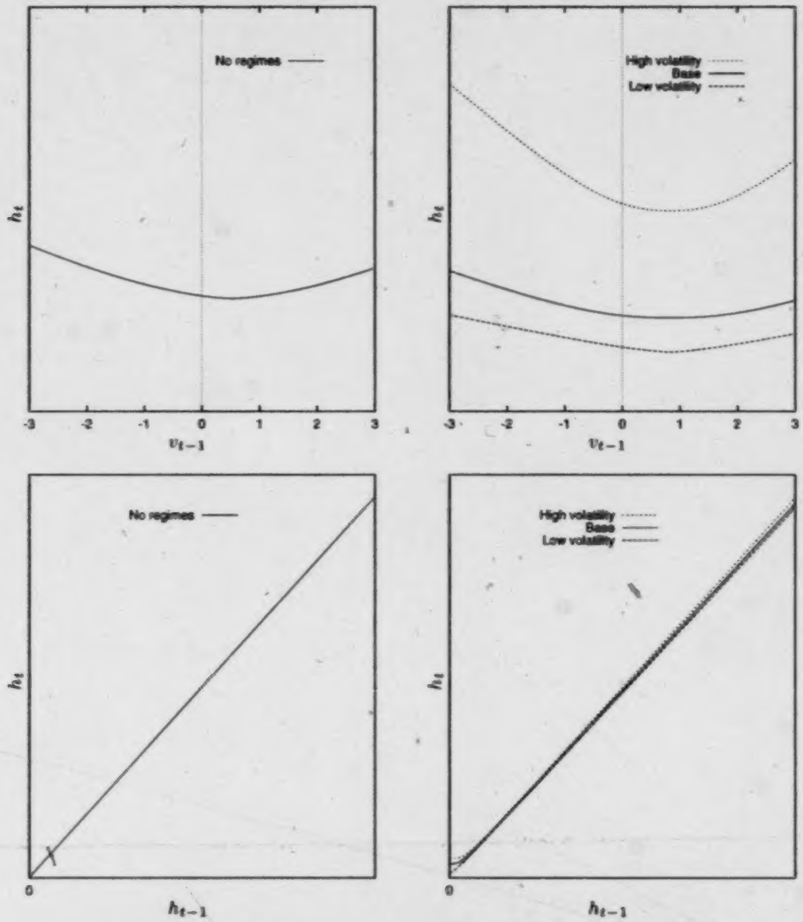

The top two panels of the figure present the NIC relating volatility shocks in the previous period to the conditional volatility in the current period for the raw return series covering the dates January 2, 1915 through December 29, 1995 (22,038 observations). The left-hand panel shows the NIC for the conventional version of the model, and the right-hand panel shows the three curves which result when the regime dependent methodology is adopted. The range of volatility shocks over which the figure is presented corresponds to the range ±3. Since $v_t \sim iid(0,1)$ this corresponds to ±3 standard deviations. Both figures include a dashed vertical line corresponding to $v_{t-1} = 0$ for comparison of asymmetry effects. The bottom two panels of the figure present the HIC relating the conditional volatility in the previous period to the conditional volatility in the current period for the raw return series covering the same period. The left-hand panel shows the HIC for the conventional version of the model, and the right-hand panel shows the three curves which result when the regime dependent methodology is adopted. To facilitate comparison, the vertical scale of each pair of figures is the same.

Figure 6.9: News and History Impact Curves 1915–1951

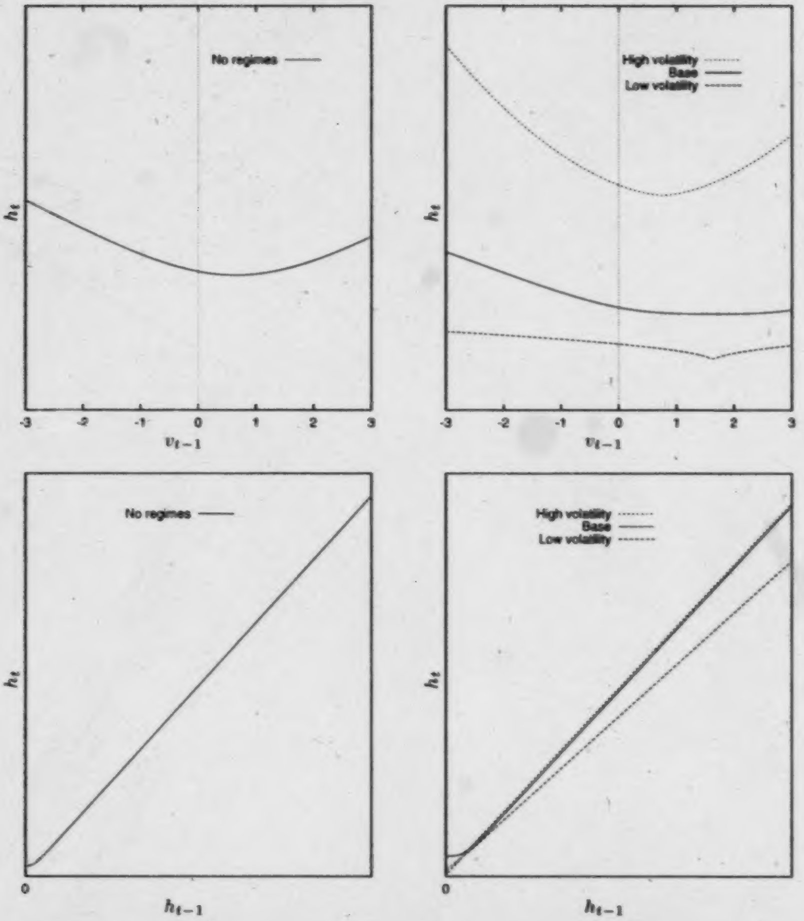

The top two panels of the figure present the NIC relating volatility shocks in the previous period to the conditional volatility in the current period for the raw return series covering the dates January 2, 1915 through December 31, 1951 (10,944 observations). The left-hand panel shows the NIC for the conventional version of the model, and the right-hand panel shows the three curves which result when the regime dependent methodology is adopted. The range of volatility shocks over which the figure is presented corresponds to the range ±3. Since $v_t \sim iid(0,1)$ this corresponds to ±3 standard deviations. Both figures include a dashed vertical line corresponding to $v_{t-1} = 0$ for comparison of asymmetry effects. The bottom two panels of the figure present the HIC relating the conditional volatility in the previous period to the conditional volatility in the current period for the raw return series covering the same period. The left-hand panel shows the HIC for the conventional version of the model, and the right-hand panel shows the three curves which result when the regime dependent methodology is adopted. To facilitate comparison, the vertical scale of each pair of figures is the same.

Figure 6.10: News and History Impact Curves 1952–1995

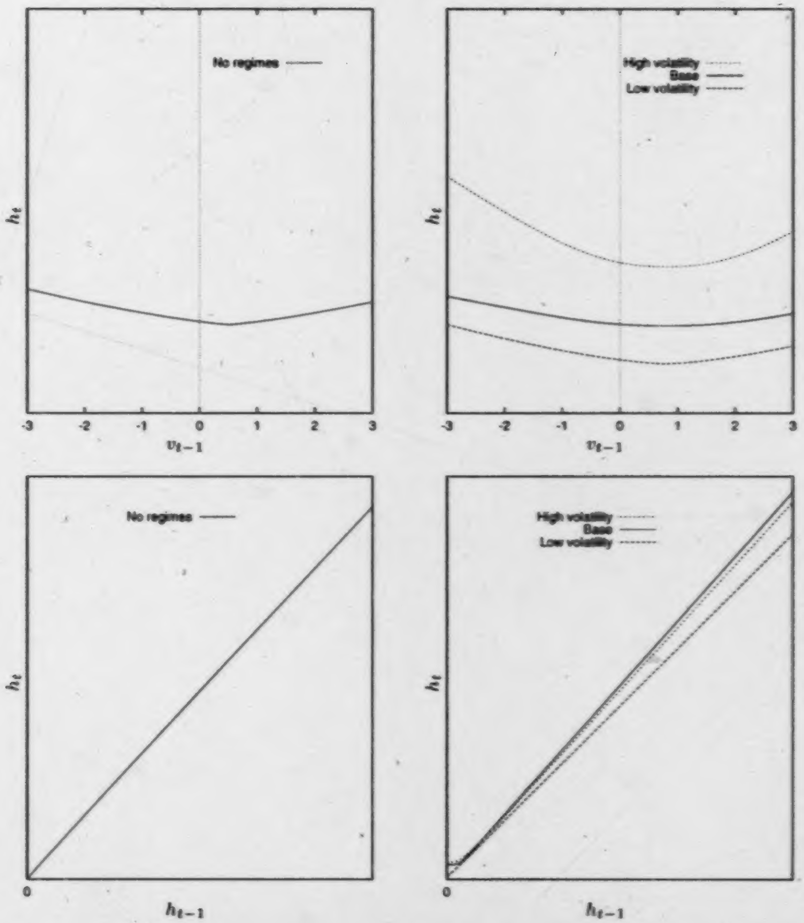

The top two panels of the figure present the NIC relating volatility shocks in the previous period to the conditional volatility in the current period for the raw return series covering the dates January 2, 1952 through December 29, 1995 (11,092 observations). The left-hand panel shows the NIC for the conventional version of the model, and the right-hand panel shows the three curves which result when the regime dependent methodology is adopted. The range of volatility shocks over which the figure is presented corresponds to the range ±3. Since $v_t \sim iid(0,1)$ this corresponds to ±3 standard deviations. Both figures include a dashed vertical line corresponding to $v_{t-1} = 0$ for comparison of asymmetry effects. The bottom two panels of the figure present the HIC relating the conditional volatility in the previous period to the conditional volatility in the current period for the raw return series covering the same period. The left-hand panel shows the HIC for the conventional version of the model, and the right-hand panel shows the three curves which result when the regime dependent methodology is adopted. To facilitate comparison, the vertical scale of each pair of figures is the same.

values for *b* presented in Table 6.12 which show an increase in value from 0.532 in the full sample to a value of 0.651 in the earlier sub-period, and a decrease to a value of 0.512 in the later sub-period. Taken together, these observations indicate that investors in the market tended to react more strongly to both positive and negative news in the earlier part of the century than in the later part. In addition, there was more asymmetry in reactions in the earlier part of the sample; while negative news causes a larger increase in subsequent conditional volatility than positive news of equal magnitude, the larger effect of negative news was more pronounced in the earlier part of the sample than in the later part. The HIC does not appear to change markedly across the three sample periods; the persistence of conditional volatility is essentially the same across the sub-periods.

Figures 6.11, 6.12, 6.13, and 6.14 present the NICs and HICs for the sub-periods 1915–1929, 1929–1951, 1952–1973, and 1974–1995 respectively. Rather than comparing these four small sub-periods to either the full sample or the two larger sub-periods, it is more revealing to compare them directly to each other. Considering first the unconditional volatility indicated by the location of the curves in space, it is apparent that the 1929–1951 sub-period exhibits the greatest unconditional volatility, and that the 1952–1973 period exhibits the least. Given the contribution of the extreme volatility associated with the Great Depression to the 1929–1952 sub-period and the low volatility associated with the 1960's to the 1952–1973 sub-period these observations are both expected and appropriate. It is interesting to note, however, that although the shape of the NIC is markedly different for the 1915–1929 and 1974–1995 sub-periods, the location in space of these two curves are similar. The unconditional volatility in the market is at a similar level in these two periods which are not characterized by large periods of unusually high or low volatility. The curvature of the NIC is most pronounced in the 1929–1951 sub-period. This suggests that during this period, investors in the market reacted more strongly to large news events, both positive and negative, than during the other periods. The earliest sub-period, 1915–1929, is characterized by a NIC which is also more curved than the periods after 1952, and in addition this curve is shifted much more to the right than any of the other periods. The parameter *b* presented in Table 6.12 reflects this large shift, taking on a value of 1.289 in the 1915–1929 sub-period, and values of 0.585, 0.515, and 0.503 in the 1929–1951, 1952–1973, and 1974–1995 sub-periods respectively. This suggests that during the period prior to the Crash of 1929, the asymmetry in the reaction of investors to negative news over positive news is stronger than during any other period—even during the period *immediately following* the Crash. This raises interesting questions regarding possible structural changes in the market coincident with the Crash. Finally, the curvature of the NIC decreases from the 1929–1951 to the 1952–1973 sub-period, and again from the 1952–1973 to the 1974–1995 sub-period.

Prior to the Crash, investors were much more concerned about negative news than about positive news. Following the Crash, the overall volatility of the market increased substantially. In addition, large news events of both signs resulted in proportionately greater subsequent volatility than small news events—not only is the market more volatile following the Crash, investors appear to react more strongly to large news events during this period. In general, since the Great Depression both the

Figure 6.11: News and History Impact Curves 1915–1929

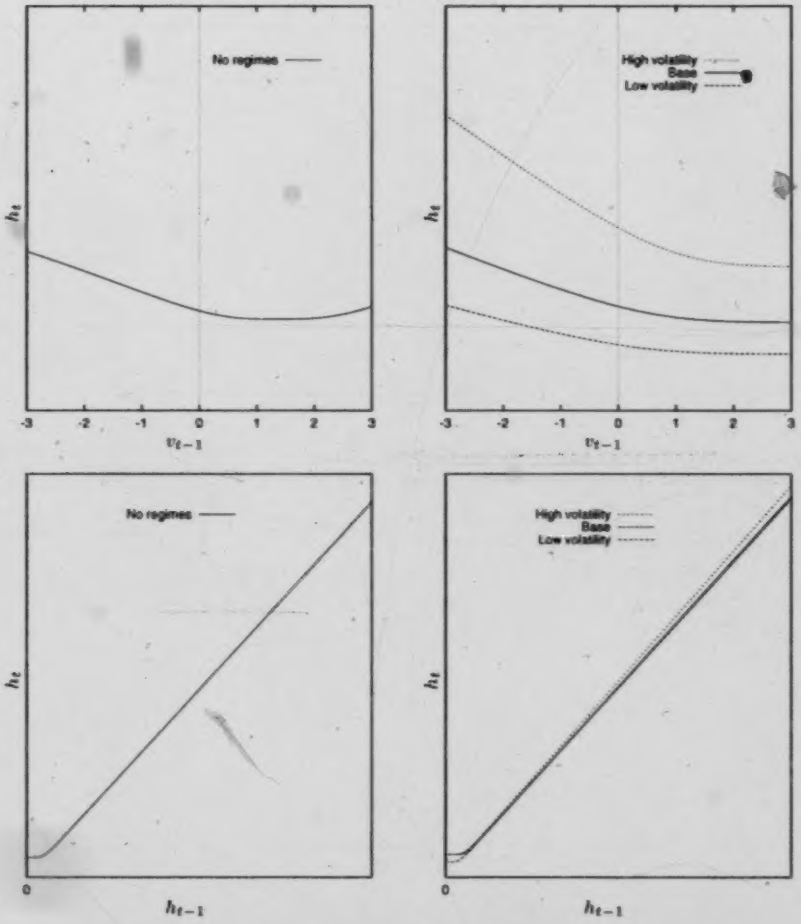

The top two panels of the figure present the NIC relating volatility shocks in the previous period to the conditional volatility in the current period for the raw return series covering the dates January 2, 1915 through September 30, 1929 (4,406 observations). The left-hand panel shows the NIC for the conventional version of the model, and the right-hand panel shows the three curves which result when the regime dependent methodology is adopted. The range of volatility shocks over which the figure is presented corresponds to the range ±3. Since $v_t \sim iid(0,1)$ this corresponds to ±3 standard deviations. Both figures include a dashed vertical line corresponding to $v_{t-1} = 0$ for comparison of asymmetry effects. The bottom two panels of the figure present the HIC relating the conditional volatility in the previous period to the conditional volatility in the current period for the raw return series covering the same period. The left-hand panel shows the HIC for the conventional version of the model, and the right-hand panel shows the three curves which result when the regime dependent methodology is adopted. To facilitate comparison, the vertical scale of each pair of figures is the same.

Figure 6.12: News and History Impact Curves 1929–1951

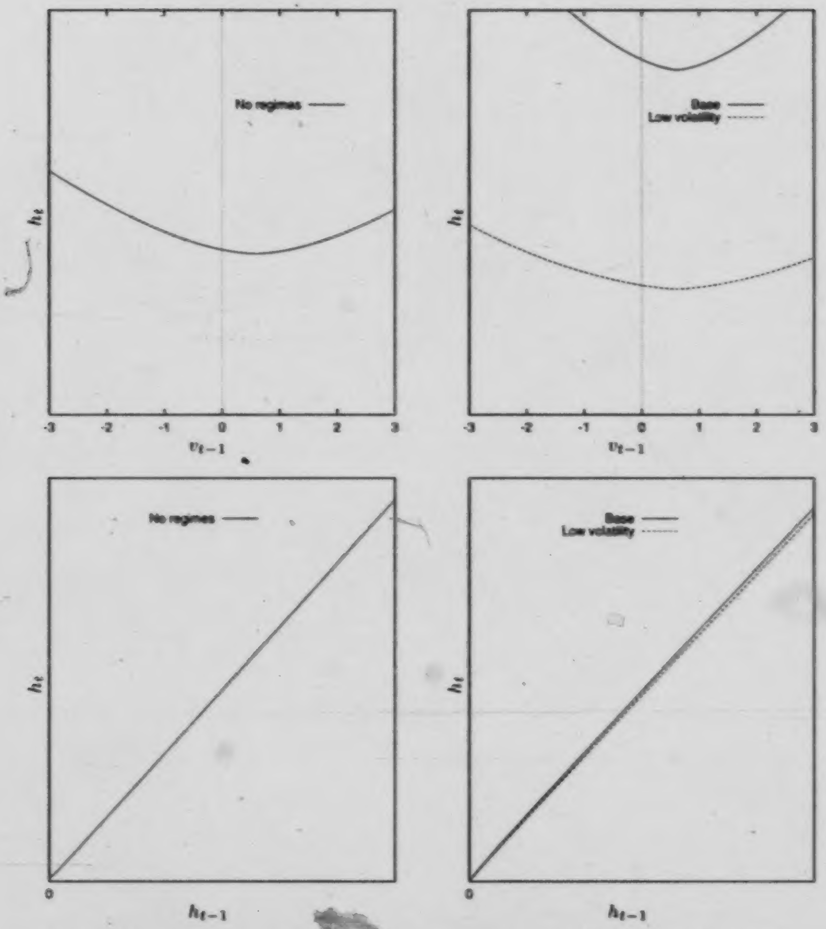

The top two panels of the figure present the NIC relating volatility shocks in the previous period to the conditional volatility in the current period for the raw return series covering the dates October 10, 1929 through December 31, 1951 (6,536 observations). The left-hand panel shows the NIC for the conventional version of the model, and the right-hand panel shows the three curves which result when the regime dependent methodology is adopted. The range of volatility shocks over which the figure is presented corresponds to the range ±3. Since $v_t \sim iid(0,1)$ this corresponds to ±3 standard deviations. Both figures include a dashed vertical line corresponding to $v_{t-1} = 0$ for comparison of asymmetry effects. The bottom two panels of the figure present the HIC relating the conditional volatility in the previous period to the conditional volatility in the current period for the raw return series covering the same period. The left-hand panel shows the HIC for the conventional version of the model, and the right-hand panel shows the three curves which result when the regime dependent methodology is adopted. To facilitate comparison, the vertical scale of each pair of figures is the same.

Figure 6.13: News and History Impact Curves 1952–1973

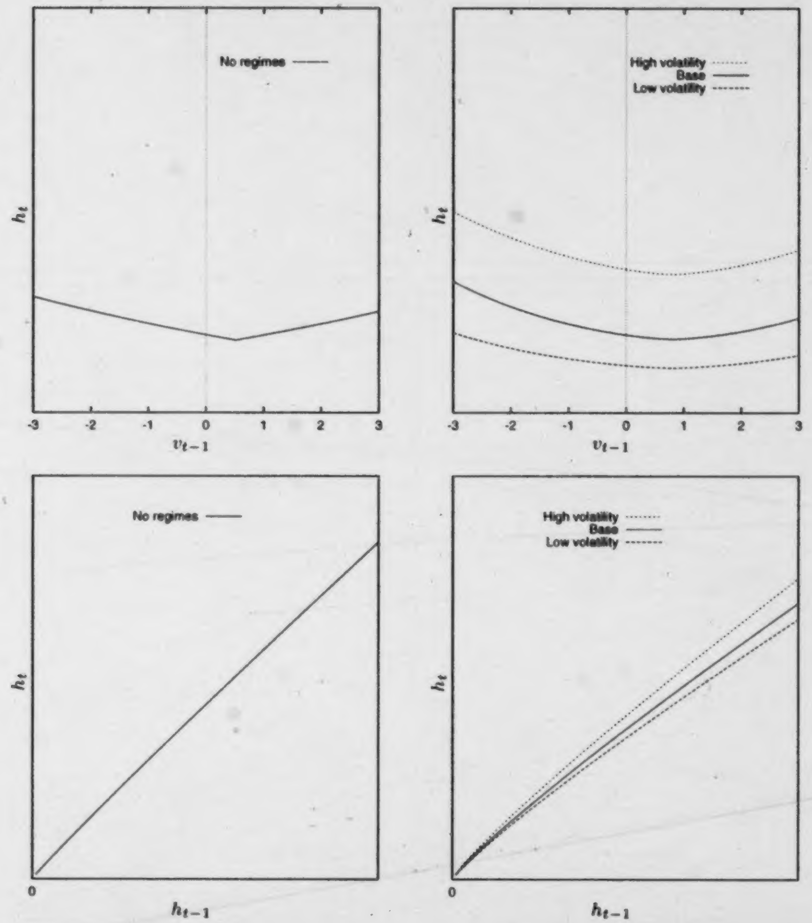

The top two panels of the figure present the NIC relating volatility shocks in the previous period to the conditional volatility in the current period for the raw return series covering the dates January 2, 1952 through December 31, 1973 (5,531 observations). The left-hand panel shows the NIC for the conventional version of the model, and the right-hand panel shows the three curves which result when the regime dependent methodology is adopted. The range of volatility shocks over which the figure is presented corresponds to the range ±3. Since $v_t \sim iid(0,1)$ this corresponds to ±3 standard deviations. Both figures include a dashed vertical line corresponding to $v_{t-1} = 0$ for comparison of asymmetry effects. The bottom two panels of the figure present the HIC relating the conditional volatility in the previous period to the conditional volatility in the current period for the raw return series covering the same period. The left-hand panel shows the HIC for the conventional version of the model, and the right-hand panel shows the three curves which result when the regime dependent methodology is adopted. To facilitate comparison, the vertical scale of each pair of figures is the same.

Figure 6.14: News and History Impact Curves 1974–1995

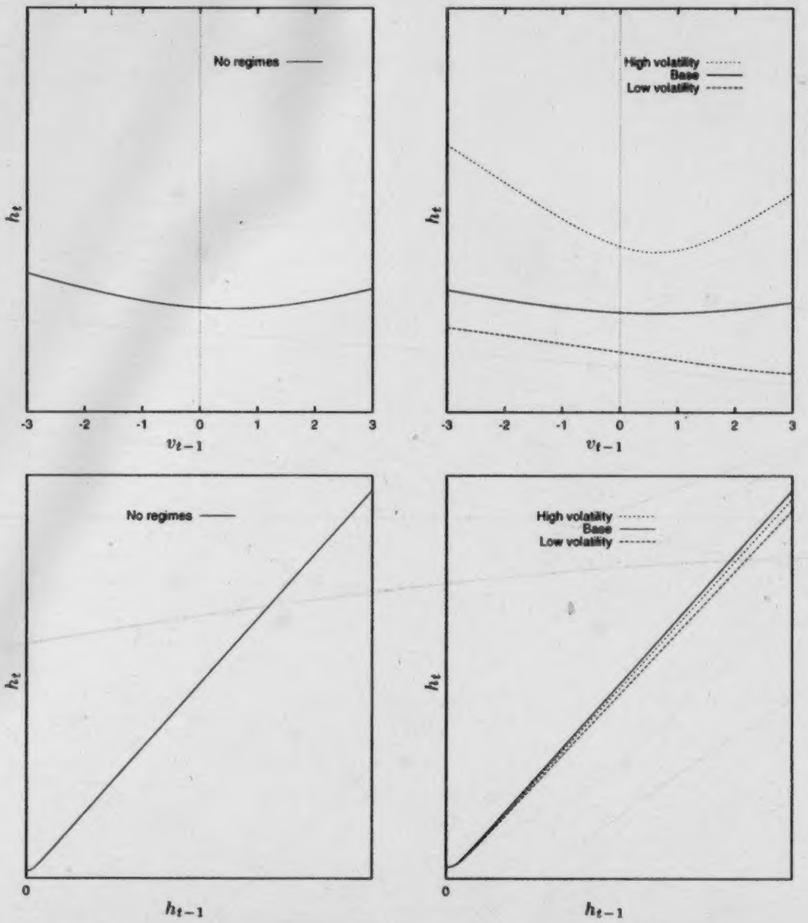

The top two panels of the figure present the NIC relating volatility shocks in the previous period to the conditional volatility in the current period for the raw return series covering the dates January 2, 1974 through December 29, 1995 (5,559 observations). The left-hand panel shows the NIC for the conventional version of the model, and the right-hand panel shows the three curves which result when the regime dependent methodology is adopted. The range of volatility shocks over which the figure is presented corresponds to the range ±3. Since $v_t \sim iid(0,1)$ this corresponds to ±3 standard deviations. Both figures include a dashed vertical line corresponding to $v_{t-1} = 0$ for comparison of asymmetry effects. The bottom two panels of the figure present the HIC relating the conditional volatility in the previous period to the conditional volatility in the current period for the raw return series covering the same period. The left-hand panel shows the HIC for the conventional version of the model, and the right-hand panel shows the three curves which result when the regime dependent methodology is adopted. To facilitate comparison, the vertical scale of each pair of figures is the same.

overall volatility of the market and the extent to which the market reacts proportionately more strongly to large news events than small news events has lessened. The overall volatility was lower during the period which includes the Korean and Vietnam wars than during the 1974–1995 period. The reverse is true for the difference in the reaction of the market to large and small news events—large news events have a proportionately greater impact on subsequent volatility in the 1952–1973 period than in the 1974–1995 period. The former period was characterized by relatively low volatility but exhibited stronger reactions to large news events, while the latter period was characterized by somewhat higher overall volatility but less aggressive reactions to large news events. One possible interpretation of the progression in reaction of conditional volatility to news presented in these four figures involves changing levels of risk aversion.

Finally, the HIC indicates that for much of the sample period there has been little difference in the persistence of conditional volatility. The exception to this was during the 1952–1973 sub-period where the NIC has a much smaller slope, indicating that the persistence during this was notably less than during any other period.

Regime Dependent Models

Table 6.13 presents the parameter estimates which result when regimes are detected and the parameter values are permitted to vary across regimes. As with Table 6.12, the first section of the table presents the results for the entire extended sample which encompasses the dates January 2, 1915 through December 29, 1995. The second section of the table presents the results for the two subsamples which cover the date ranges January 2, 1915 through December 31, 1951 and January 2, 1952 through December 29, 1995. The final section of the table presents the results for the four subsamples covering the date ranges January 2, 1915 through September 30, 1929, October 1, 1929 through December 31, 1951, January 2, 1952 through December 31, 1973, and January 2, 1974 through December 29, 1995.

As was the case in Section 5.4, comparison of the mean equation parameters, γ_0, γ_2, γ_3, and γ_4, between the conventional and regime dependent versions of the model presented in Tables 6.12 and 6.13 respectively indicates that almost no difference results from allowing the conditional variance to depend upon the regime. In all cases the parameter values are similar for the conventional and regime dependent versions of the model. While numerical differences do exist, the differences are small and in no case is there a change in the significance of the parameter estimates.

The conditional volatility equation parameters exhibit considerably more activity than the mean equation parameters. Following the procedure employed above, the discussion of the conditional volatility equation parameters will be conducted with reference to the NICs and HICs rather than the tabulated parameter estimates because of the difficulties involved in interpreting the interrelationships between the various parameters. The single exception to this is parameter b which controls the asymmetry in the NIC by allowing it to shift. As the asymmetry in response controlled by this parameter is independent of the other parameters of the equation, it is meaningful to interpret the parameter by itself. In each case, the top right panel of the relevant figure presents the NIC for the regime dependent version of the model, and the bottom right panel presents the corresponding HIC.

The NICs and HICs for the regime dependent version of the model for the 1926–1990 sub-period are presented in Figure 6.4, and the corresponding curves for the 1915–1995 sample are presented in Figure 6.8. The NICs are similar, with the unconditional volatility of the high volatility regimes being slightly greater in the shorter sample, as indicated by the slightly higher location in space of the relevant curve. There is also a slight shift to the right of the curves for both the low and base volatility regimes in the longer sample, indicating that there is slightly more asymmetry in the response of conditional volatility to news in the longer period. This result is not surprising given the large degree of asymmetry shown in the conventional model for the 1915–1929 sample period. The location in space and curvature of the curves for the base and low volatility regimes exhibit no differences between the sample periods. The HICs are also similar between the two sample periods. The HIC for the 1915–1995 sample period indicates that there is no difference in the degree of persistence between the three regime types. The persistence exhibited by the high and base volatility regimes for the 1926–1990 sample period are equal, and are the same as the persistence for all three regime types in the longer sample. Low volatility regimes exhibit slightly less persistence in conditional volatility for the smaller sample period.

Comparing the full 1915–1995 sample period presented in Figure 6.8 with the 1915–1951 sub-period presented in Figure 6.9 reveals several notable differences. The location in space of the NIC indicates that the unconditional volatility of high volatility regimes in the 1915–1929 sub-period is larger than the unconditional volatility of high volatility regimes in the full sample, although the curvature is similar. Again, the extreme volatility of the 1930's is reflected in the unconditional volatility of high volatility regimes, although interestingly the nature of the impact of news on subsequent conditional volatility is essentially the same for high volatility regimes in both sample periods. The location in space, and hence the unconditional volatility, of the NIC for the base and low volatility regimes is approximately equal for both the sample periods. In addition, the curvature of the NIC is similar for the base regimes. The curve is shifted to the right for the shorter sample period, repeating the earlier observation that the asymmetry in response of positive versus negative news is greater in the earlier part of the sample period. The tabulated results for the parameter b also illustrate this effect, and in addition indicate that the asymmetry is more pronounced for base and low volatility regimes ($b = 1.637$) than it is for high volatility regimes ($b = 1{,}637 - 0.866 = 0.771$). The comparable figures for the longer sample period are $b = 0.828$ for all three regime types, indicating that the asymmetry difference is largely restricted to the base and low volatility regimes; there does not appear to be a large difference in the asymmetry of the impact of news on conditional volatility for high volatility regimes across the two sample periods. Finally, the curvature of the NIC for low volatility regimes is much less for the smaller sample period than it is for the larger period. Comparison of the HICs between the two sample periods reflects a similar situation to that observed in the comparison of the 1915–1995 sample period with the 1926–1990 period. The persistence in conditional volatility exhibited by high and base volatility regimes in the 1915–1951 sample period are equal, and is the same as the persistence for all three regime types in the full 1915–1995 sample. Low volatility regimes exhibit less persistence in conditional volatility for the smaller sample period, and this difference is larger than was the difference between the persistence in the low volatility regimes for the 1915–1995 and 1926–1990 periods.

Comparing the full 1915–1995 sample period with the 1952–1995 sub-period presented in Figure 6.10 also reveals notable differences. The location in space of the NIC indicates that the unconditional volatility of all three regime types is less in the later sample period than the full sample, which is expected given the absence of the high volatility period of the 1930's and the presence of the low volatility period of the 1960's. Other than the difference in unconditional volatility, the differences between the NICs resulting from regime dependence are unremarkable. The curvature of the curves for low volatility regimes are nearly identical, and the curves for the base regimes exhibit slightly less curvature in the later sub-period. The difference for the high volatility regimes is larger, with the later sub-period again exhibiting less curvature. Taken together, these observations tend to reinforce the observation made earlier that there has been a shift in the reaction of the market to news, with the impact of news on subsequent volatility being less in the later part of the sample period. There is more divergence in the persistence of conditional volatility for the 1952–1995 sample period than for the 1915–1951 or 1915–1995 periods. In the later sub-period, the persistence of conditional volatility in high volatility regimes is nearly identical to the persistence for all regime types in the full sample. The base regime exhibits slightly more persistence, and the low volatility regimes exhibit somewhat less persistence.

Comparison of the NICs and HICs for the conventional versions of the model with the regime dependent versions for the for small sub-periods indicates that the general observations made above regarding the conventional versions are also applicable to the regime dependent versions. For this reason, it is appropriate to consider the impact of allowing for regime dependence in the model for each of the small sub-periods individually.

Figure 6.11 presents the NICs and HICs for the 1915–1929 sub-period. The separation of the three NICs resulting from the regime dependent version of the model reflects the difference in unconditional volatility present in the three regimes. The asymmetry noted in the conventional version of the model is also present all three NICs in the regime dependent version, with the shift appearing to be greater in the regime dependent version of the model than in the conventional version. It is similar for all three regime types. There is slightly greater curvature present in the base regime than in low volatility regimes, and somewhat more curvature in high volatility regimes than in the base regime. This indicates that during this period higher unconditional volatility in the market was associated with an increased impact of news on subsequent conditional volatility and, as the curves are all shifted dramatically to the right, the impact was largely restricted to negative news events. The differences between HICs for the conventional and regime dependent versions of the model are unremarkable. There is a slight increase in the persistence of conditional volatility in high volatility regimes, but no change in the base or low volatility regimes associated with allowing parameters to vary across regime types.

The NICs and HICs for the 1929–1951 sub-period, presented in Figure 6.12, are different from the other three sub-periods in that there were only base and low volatility regimes detected in this sample period. As outlined earlier, the reason for this stems from the dominance of the period of extremely high volatility associated with the Great Depression. The magnitude of this effect is illustrated in Figure 6.6. Because the high volatility period dominates the sample period, it is detected as belonging to the base regime rather than a high volatility regime. For this reason, it would be more appropriate to consider the curves labeled "Base Regime" in the two right panels of Figure 6.12 as representing high volatility regimes. Similarly, the curves labeled "Low Volatility" should be interpreted as representing the base regime. Unfortunately no true low volatility regimes have been detected and isolated. The curvature of the NIC for high volatility regimes in the regime dependent version of the model is much greater than the curvature of the NIC for the base regime. As the major portion of the high volatility regimes are in the early part of the sample period,[3] the greater curvature of the NIC observed in high volatility regimes lends additional support to the observation made earlier that the impact of news on subsequent conditional volatility has been decreasing through time. The HIC shows a small amount of activity when regime dependence is permitted, with the persistence in both high and base volatility regimes decreasing from the level which resulted from the conventional model. In addition, the base regime exhibits slightly less persistence than is exhibited by high volatility regimes.

The NICs for the regime dependent version of the model for the 1952–1973 sub-period is presented in Figure 6.13, and is particularly unremarkable. The main difference between the regime types is reflected in the separation of the three NICs, indicating that the unconditional volatility varies among the regime types, but that the impact of news on subsequent conditional volatility is essentially the same for the three types of regimes. There is a slight difference in the curvature exhibited by the three curves, with high and base volatility regimes appearing to have the same level of curvature, and the NIC for low volatility regimes exhibiting slightly less. The curvature of the NIC for high and base volatility regimes is somewhat greater than that exhibited by the NIC for the conventional version of the model, and the curvature of the NIC for low volatility regimes is slightly less than that exhibited by the conventional model. The HICs are more interesting. The level of persistence in conditional volatility is substantially less for all three regime types than it was for the conventional version of the model. In addition, there are considerable differences in the level of persistence exhibited by the three regime types. The high volatility regimes exhibit the greatest level of persistence, the low volatility regimes exhibit the least, and the base regime is between the two extremes.

In contrast to the unremarkability of the NICs for the regime dependent version of the model for the 1952–1973 sub-period, the curves for the 1974–1995 sub-period, presented in Figure 6.14, exhibit the greatest differences between the three regime types of any sub-period considered. The curve for the base regime is similar to the curve obtained with the conventional version of the model, both in curvature and location in space, although the regime dependent version does exhibit slightly less curvature than its conventional counterpart. Removal of the periods of high and low volatility does not seem to have a major impact on the response of conditional volatility to news.

---

[3] 906 of 1,309 days, or 69% of high volatility regimes occur prior to April 24, 1935 which represents the first quarter of the sample period. In addition, 1,307 of the 1,309 days of high volatility regimes occur prior to November 15, 1940 which represents the first half of the sample period.

In addition to being located higher in space, the NIC for high volatility regimes also exhibits the second greatest curvature of any NIC from any of the sub-periods examined. The only period where the difference in impact between large and small news events, both positive and negative, was greater than during the high volatility periods in the 1974–1995 sub-period was during the 1929–1952 sub-period. In addition, since the high volatility regimes associated with the 1929–1952 sub-period are restricted almost entirely to the first half of that period, the only period which exhibited a NIC with a greater curvature than that exhibited during high volatility regimes of the 1974–1995 sub-period is in fact the period immediately following the Crash of 1929, encompassing the Great Depression and ending just over a year prior to the entry of the United States into World War II.

The NIC for low volatility regimes is also remarkable. The curve exhibits moderate curvature, but is shifted dramatically to the right. The parameter b for high and base volatility regimes takes on a value of 0.641, while for low volatility regimes the value of b is $0.641 + 2.778 = 3.419$, a shift that is nearly twice as great as that observed in the regime dependent version of the model during the 1915–1929 sub-period. Also of interest is the fact that all low volatility regimes occur in the later portion of the sub-period,[4] and the sub-period ends during a low volatility regime. This may have an impact on the nature of the volatility forecasts produced by the model and discussed in Chapter 7.

The HIC for the regime dependent version of the model shows a slight difference in the degree of persistence of conditional volatility among the three regime types. As was the case with the NIC, the HIC for the base regime is similar to the HIC which resulted from the conventional version of the model. Both high and low volatility regimes exhibit slightly less persistence, with high volatility regimes exhibiting more persistence than low volatility regimes, but less than the base regime.

Summary

Initial examination of the NICs and HICs for the conventional version of the model would seem to indicate that, with the exception of the extreme volatility period associated with the 1930's, the nature of the impact of news on subsequent conditional volatility is reasonably stable across sample periods. While there are differences between the NICs in terms of location in space, curvature, and asymmetry (shift), the curves for the overall 1915–1995 period are grossly similar to those for the 1915–1951 and 1952–1995 sub-periods, and the curves for the full period and larger sub-periods are grossly similar to those for the 1915–1929, 1929–1951, 1952–1973, and 1974–1995 sub-periods. In particular, the results for the 1951–1995 sub-period are quite similar to those for both the 1951–1973 and 1974–1995 sub-periods.

---

[4] The first low volatility regime begins on October 5, 1992 which is 6,851 calendar days after the sub-period begins, or 85% of the way through the period.

The results from the application of the RD-GARCH model, however, are interesting in their diversity. In particular, the curves from the regime dependent model applied to the 1915–1929, 1929–1951, 1952–1973, and 1974–1995 sub-periods are all notably different from each other, and in particular the 1974–1995 sub-period is dramatically different from the 1952–1973 sub-period. While these differences may prove useful in understanding historical changes in volatility, it is their potential usefulness in generating improved forecasts of volatility which may be of more immediate economic benefit. This is the subject that will be investigated in Chapter 7.

Table 6.1: Dates of regimes 1926–1990 (raw returns)

Base		Low Volatility				High Volatility			
$\bar{r}$	Var	From	To	$\bar{r}$	Var	From	To	$\bar{r}$	Var
4.537×10^{-4}	6.068×10^{-5}	1926 11 05	1927 01 26	2.448×10^{-4}	2.206×10^{-5}	1929 01 09	1938 10 25	-4.388×10^{-5}	4.055×10^{-4}
9796 Days		1927 03 23	1927 06 04	1.278×10^{-3}	2.218×10^{-5}	1938 12 05	1939 02 28	3.299×10^{-4}	1.797×10^{-4}
		1944 01 06	1944 03 24	3.591×10^{-4}	2.117×10^{-5}	1939 03 09	1939 12 11	-8.361×10^{-5}	2.249×10^{-4}
		1944 04 19	1944 07 19	1.499×10^{-3}	2.098×10^{-5}	1940 02 26	1940 10 07	-4.698×10^{-4}	2.055×10^{-4}
		1950 01 13	1950 04 04	8.760×10^{-4}	2.089×10^{-5}	1946 07 10	1947 02 17	-7.267×10^{-4}	1.986×10^{-4}
		1951 07 26	1951 10 20	9.235×10^{-4}	2.341×10^{-5}	1973 10 22	1974 02 06	-2.215×10^{-3}	1.806×10^{-4}
		1951 11 08	1952 01 31	1.309×10^{-3}	2.084×10^{-5}	1974 05 22	1975 04 30	9.014×10^{-5}	1.897×10^{-4}
		1952 02 20	1952 10 16	3.499×10^{-4}	2.187×10^{-5}	1982 07 23	1982 11 12	3.388×10^{-3}	1.839×10^{-4}
		1952 11 03	1953 02 09	1.056×10^{-3}	2.202×10^{-5}	1987 02 23	1988 10 05	-2.026×10^{-5}	2.616×10^{-4}
		1953 10 30	1954 04 30	1.349×10^{-3}	2.096×10^{-5}	4384 Days			
		1957 02 14	1957 07 30	1.092×10^{-3}	2.039×10^{-5}				
		1961 01 05	1961 04 04	2.041×10^{-3}	2.033×10^{-5}				
		1961 09 26	1962 04 10	3.977×10^{-4}	1.970×10^{-5}				
		1962 05 29	1970 06 01	2.986×10^{-4}	4.229×10^{-5}				
		1971 01 05	1971 05 18	1.343×10^{-3}	2.171×10^{-5}				
		1972 01 25	1972 04 26	9.224×10^{-4}	2.058×10^{-5}				
		3304 Days							

The table presents the dates, mean return, and unconditional variance of the return for each of the regimes identified in the raw return series. The base regime covers the entire sample period from January 2, 1926 through December 31, 1990 (17,484 observations), with the returns for high and low volatility regimes excluded from the calculation of the sample statistics. Each row in the sections of the table labeled "Low Volatility" and "High Volatility" represent a single period identified as either a regime of unusually low or unusually high unconditional volatility respectively. The number of days identified at the end of each column identifies the total number of days identified with that particular type of regime in the sample period.

Table 6.2: Dates of regimes 1926–1990 (excess returns)

Base		Low Volatility				High Volatility			
$\bar{r}$	Var	From	To	$\bar{r}$	Var	From	To	$\bar{r}$	Var
3.531×10^{-4}	6.089×10^{-5}	1926 11 09	1927 01 26	2.434×10^{-5}	2.194×10^{-5}	1929 03 21	1939 01 17	-1.103×10^{-4}	4.052×10^{-4}
		1927 03 23	1927 06 04	1.186×10^{-3}	2.218×10^{-5}	1939 03 14	1939 12 11	-3.583×10^{-5}	2.281×10^{-4}
		1944 01 06	1944 03 24	3.503×10^{-4}	2.117×10^{-5}	1940 02 23	1940 10 07	-5.290×10^{-4}	2.037×10^{-4}
		1944 04 19	1944 07 19	1.490×10^{-3}	2.098×10^{-5}	1946 07 09	1947 02 17	-6.991×10^{-4}	1.977×10^{-4}
		1950 01 13	1950 03 31	6.748×10^{-4}	2.015×10^{-5}	1973 10 22	1974 02 06	-2.411×10^{-3}	1.806×10^{-4}
		1951 07 25	1951 10 20	7.833×10^{-4}	2.365×10^{-5}	1974 05 22	1975 05 08	7.461×10^{-5}	1.884×10^{-4}
		1951 11 08	1952 01 31	1.269×10^{-3}	2.084×10^{-5}	1982 07 20	1982 11 12	3.154×10^{-3}	1.775×10^{-4}
		1952 02 20	1952 10 16	3.062×10^{-4}	2.186×10^{-5}	1986 08 25	1988 04 28	7.283×10^{-5}	2.617×10^{-4}
		1952 11 03	1953 02 09	1.006×10^{-3}	2.202×10^{-5}				
		1953 10 30	1954 04 30	1.322×10^{-3}	2.096×10^{-5}				
		1957 02 14	1957 07 30	1.012×10^{-3}	2.038×10^{-5}				
		1961 01 05	1961 04 04	1.981×10^{-3}	2.034×10^{-5}				
		1961 09 26	1962 04 05	3.625×10^{-4}	1.940×10^{-5}				
		1962 05 29	1970 06 01	1.828×10^{-4}	4.232×10^{-5}				
		1971 01 28	1971 05 18	8.945×10^{-4}	2.235×10^{-5}				
		1972 01 25	1972 04 26	8.345×10^{-4}	2.059×10^{-5}				
9860 Days		3279 Days				4345 Days			

The table presents the dates, mean excess return, and unconditional variance of the excess return for each of the regimes identified in the excess return series. The base regime covers the entire sample period from January 2, 1926 through December 29, 1990 (17,484 observations), with the returns for the high and low volatility regimes excluded from the calculation of the sample statistics. Each row in the sections of the table labeled "Low Volatility" and "High Volatility" represent a single period identified as either a regime of unusually low or unusually high unconditional volatility respectively. The number of days identified at the end of each column identifies the total number of days identified with that particular type of regime in the sample period. This table is a duplicate of Table 5.1, and is presented here to facilitate comparison with Table 6.1.

Table 6.3: Estimates of model parameters (raw vs. excess returns, no regimes)

Model	γ_0	γ_2	γ_3	γ_4	λ	ν	b	ω	α	β
1926–1990 (excess)	$5.856^{[-4]}$ $(5.780^{[-5]})^*$	0.165 $(0.009)^*$	−0.035 $(0.008)^*$	−0.002 $(1.437^{[-4]})^*$	0.941 $(0.265)^*$	1.538 $(0.114)^*$	0.489 $(0.056)^*$	$1.861^{[-4]}$ $(2.364^{[-4]})$	0.059 $(0.016)^*$	0.926 $(0.018)^*$
1926–1990 (raw)	$6.647^{[-4]}$ $(5.801^{[-5]})^*$	0.165 $(0.009)^*$	−0.036 $(0.008)^*$	−0.002 $(1.453^{[-4]})^*$	0.923 $(0.252)^*$	1.527 $(0.111)^*$	0.493 $(0.056)^*$	$2.028^{[-4]}$ $(2.441^{[-4]})$	0.059 $(0.016)^*$	0.927 $(0.018)^*$

Entries of the form $x^{[-y]}$ represent scientific notation; that is $2.5^{[-2]}$ represents the value 2.5×10^{-2}.

The table presents the subsample parameter estimates for the conventional GARCH model. Numbers in parentheses under the parameter estimates represent asymptotic standard errors. Parameters which are significantly different from zero at the 5% level are indicated with an asterisk. The data set comprises returns on an U.S. equity index covering the period January 2, 1926 to December 29, 1990, a total of 17,484 observations.

Table 6.4: Estimates of model parameters (raw vs. excess returns)

Model	γ_0	γ_2	γ_3	γ_4	λ	ν	b	ω	α	β
1926–1990	$6.311^{[-4]}$	0.162	−0.034	−0.002	4.721	2.587	0.754	$5.890^{[-12]}$	0.096	0.741
(excess)	($5.597^{[-5]}$)*	(0.009)*	(0.008)*	($1.405^{[-4]}$)*	(1.652)*	(0.361)*	(0.051)*	($4.785^{[-11]}$)	(0.031)*	(0.073)*
					−3.794	−1.333	0	$4.407^{[-4]}$	0	0.094
					(1.527)*	(0.402)*		($8.079^{[-4]}$)		(0.052)
					−3.420	−0.905	0	$1.307^{[-4]}$	0	0.106
					(1.484)*	(0.442)*		($3.327^{[-4]}$)		(0.050)*
1926–1990	$7.147^{[-4]}$	0.163	−0.035	−0.002	2.854	2.014	0.725	$6.400^{[-8]}$	0.096	0.787
(raw)	($5.635^{[-5]}$)*	(0.009)*	(0.008)*	($1.441^{[-4]}$)*	(1.504)	(0.268)*	(0.052)*	($4.791^{[-7]}$)	(0.034)*	(0.060)*
					−1.529	−0.752	0	$5.446^{[-5]}$	0.049	0
					(1.202)	(0.321)*		($1.536^{[-4]}$)	(0.033)	
					−1.104	0	0	$1.745^{[-5]}$	0	0.048
					(0.805)			($6.212^{[-5]}$)		(0.024)*

Entries of the form $x^{[-y]}$ represent scientific notation; that is $2.5^{[-2]}$ represents the value 2.5×10^{-2}.

The table presents the subsample parameter estimates for the Regime Dependent GARCH model. Numbers in parentheses under the parameter estimates represent asymptotic standard errors. Parameters which are significantly different from zero at the 5% level are indicated with an asterisk. The first pair of rows in each section present the parameter estimates and standard errors for the base regime, and the remaining four rows present the parameter estimates and standard errors for the regimes of unusually low and high volatility respectively. Parameters which are excluded by the stepwise estimation procedure are represented by a value of zero, with no corresponding standard error presented. The data set comprises returns on an U.S. equity index covering the period January 2, 1926 to December 29, 1990, a total of 17,484 observations.

Table 6.5: Dates of regimes 1915–1995

Base		Low Volatility				High Volatility			
$\bar{r}$	Var	From	To	$\bar{r}$	Var	From	To	$\bar{r}$	Var
4.667×10^{-4}	5.862×10^{-5}	1918 06 13	1918 08 24	7.166×10^{-3}	2.005×10^{-4}	1916 12 05	1917 02 26	-1.481×10^{-2}	1.719×10^{-3}
		1923 01 08	1923 03 24	9.955×10^{-3}	1.952×10^{-4}	1917 09 28	1918 03 07	-2.210×10^{-3}	1.771×10^{-3}
		1924 06 12	1924 09 05	1.421×10^{-2}	2.037×10^{-4}	1928 11 19	1939 05 09	-7.751×10^{-4}	3.929×10^{-3}
		1925 06 16	1925 09 02	1.114×10^{-2}	2.064×10^{-4}	1939 05 18	1939 12 11	9.279×10^{-3}	1.897×10^{-3}
		1926 11 11	1927 01 26	2.528×10^{-3}	2.171×10^{-4}	1940 02 14	1940 10 07	-5.100×10^{-3}	1.967×10^{-3}
		1927 03 23	1927 06 04	1.278×10^{-2}	2.218×10^{-4}	1946 07 03	1947 02 19	-7.083×10^{-3}	1.922×10^{-3}
		1944 04 19	1944 07 18	1.641×10^{-2}	1.973×10^{-4}	1955 09 01	1955 11 30	1.082×10^{-2}	1.758×10^{-3}
		1950 01 13	1950 03 30	8.411×10^{-3}	1.929×10^{-4}	1962 03 21	1962 06 18	-3.723×10^{-2}	2.318×10^{-3}
		1951 07 27	1951 10 20	8.278×10^{-3}	2.323×10^{-4}	1973 10 18	1974 02 06	-2.130×10^{-2}	1.761×10^{-3}
		1951 11 08	1952 01 31	1.309×10^{-2}	2.084×10^{-4}	1974 05 22	1975 05 09	2.755×10^{-3}	1.876×10^{-3}
		1952 02 20	1952 10 16	3.499×10^{-3}	2.187×10^{-4}	1982 07 16	1982 11 12	3.305×10^{-2}	1.736×10^{-3}
		1952 11 03	1953 02 06	1.270×10^{-2}	1.937×10^{-4}	1986 12 23	1988 10 26	3.684×10^{-3}	2.397×10^{-3}
		1953 11 12	1954 04 30	1.474×10^{-2}	1.999×10^{-4}				
		1957 02 14	1957 07 30	1.092×10^{-2}	2.039×10^{-4}				
		1961 09 26	1962 01 08	9.460×10^{-3}	1.977×10^{-4}				
		1962 06 19	1970 05 28	2.712×10^{-3}	3.858×10^{-4}				
		1971 01 28	1971 05 07	1.427×10^{-2}	1.924×10^{-4}				
		1972 01 28	1972 04 26	7.883×10^{-3}	1.985×10^{-4}				
		1992 10 05	1993 02 17	7.932×10^{-3}	2.578×10^{-4}				
		1993 03 09	1994 02 07	2.814×10^{-3}	2.157×10^{-4}				
		1994 05 19	1994 08 29	6.741×10^{-3}	1.988×10^{-4}				
		1994 11 22	1995 12 29	1.085×10^{-2}	1.999×10^{-4}				
13189 Days		3993 Days				4857 Days			

The table presents the dates, mean return, and unconditional variance of the return for each of the regimes identified in the raw return series. The base regime covers the entire sample period from January 2, 1915 through December 29, 1995 (22,038 observations), with the returns for high and low volatility regimes excluded from the calculation of the sample statistics. Each row in the sections of the table labeled "Low Volatility" and "High Volatility" represent a single period identified as either a regime of unusually low or unusually high unconditional volatility respectively. The number of days identified at the end of each column identifies the total number of days identified with that particular type of regime in the sample period.

Table 6.6: Dates of regimes 1915–1951

Base		Low Volatility				High Volatility			
$\bar{r}$	Var	From	To	$\bar{r}$	Var	From	To	$\bar{r}$	Var
4.353×10^{-4}	6.832×10^{-5}	1918 06 05	1918 09 07	6.324×10^{-3}	2.207×10^{-4}	1917 09 28	1918 02 20	-8.401×10^{-4}	1.916×10^{-3}
		1922 01 12	1922 03 30	1.578×10^{-2}	2.287×10^{-4}	1929 03 13	1930 02 15	-1.180×10^{-3}	4.877×10^{-3}
		1922 12 29	1923 03 27	7.211×10^{-3}	2.178×10^{-4}	1930 02 18	1936 01 14	-1.342×10^{-3}	4.639×10^{-3}
		1923 11 20	1924 02 14	1.002×10^{-2}	2.310×10^{-4}	1936 02 10	1938 10 25	7.086×10^{-4}	2.809×10^{-3}
		1924 06 12	1924 09 06	1.243×10^{-2}	2.235×10^{-4}	1939 03 02	1939 12 04	1.731×10^{-4}	2.261×10^{-3}
		1925 06 09	1925 09 03	9.420×10^{-3}	2.392×10^{-4}	1940 03 01	1940 09 16	-7.735×10^{-3}	2.237×10^{-3}
		1926 10 28	1927 01 26	3.088×10^{-3}	2.238×10^{-4}	1946 08 07	1947 02 17	-6.156×10^{-3}	2.143×10^{-3}
		1927 03 23	1927 06 04	1.278×10^{-2}	2.218×10^{-4}				
		1944 01 05	1944 03 24	5.204×10^{-3}	2.254×10^{-4}				
		1944 04 19	1944 07 19	1.499×10^{-2}	2.098×10^{-4}				
		1944 09 14	1944 12 27	7.816×10^{-3}	2.219×10^{-4}				
		1945 04 17	1945 06 29	7.578×10^{-3}	2.861×10^{-4}				
		1950 01 13	1950 04 13	1.113×10^{-2}	2.175×10^{-4}				
		1951 07 25	1951 10 20	8.248×10^{-3}	2.365×10^{-4}				
6432 Days		1003 Days				3509 Days			

The table presents the dates, mean return, and unconditional variance of the return for each of the regimes identified in the raw return series. The base regime covers the entire sample period from January 2, 1915 through December 31, 1951 (10,944 observations), with the returns for high and low volatility regimes excluded from the calculation of the sample statistics. Each row in the sections of the table labeled "Low Volatility" and "High Volatility" represent a single period identified as either a regime of unusually low or unusually high unconditional volatility respectively. The number of days identified at the end of each column identifies the total number of days identified with that particular type of regime in the sample period.

Table 6.7: Dates of regimes 1952-1995

Base		Low Volatility				High Volatility			
$\bar{r}$	Var	From	To	$\bar{r}$	Var	From	To	$\bar{r}$	Var
4.768×10^{-4}	4.968×10^{-5}	1952 04 15	1952 10 16	2.258×10^{-4}	1.989×10^{-5}	1955 08 26	1955 12 05	1.088×10^{-3}	1.587×10^{-4}
7451 Days		1953 12 23	1954 04 01	1.446×10^{-3}	1.781×10^{-5}	1957 09 05	1957 12 04	-1.184×10^{-3}	1.507×10^{-4}
		1957 02 18	1957 07 05	1.295×10^{-3}	1.901×10^{-5}	1962 03 23	1962 09 12	-1.397×10^{-3}	1.650×10^{-4}
		1961 09 26	1962 01 08	9.460×10^{-4}	1.977×10^{-5}	1970 03 06	1970 07 28	-1.768×10^{-3}	1.640×10^{-4}
		1962 11 13	1966 07 26	5.022×10^{-4}	2.116×10^{-5}	1973 10 18	1974 02 15	-2.194×10^{-3}	1.654×10^{-4}
		1967 06 07	1967 11 21	3.365×10^{-4}	1.823×10^{-5}	1974 05 15	1975 06 19	2.021×10^{-4}	1.730×10^{-4}
		1968 07 23	1969 01 07	6.310×10^{-4}	1.934×10^{-5}	1980 01 28	1980 04 30	-1.015×10^{-3}	1.568×10^{-4}
		1992 10 08	1993 02 17	9.625×10^{-4}	2.538×10^{-5}	1982 07 07	1982 11 12	3.335×10^{-3}	1.619×10^{-4}
		1993 05 20	1994 02 03	5.456×10^{-4}	1.651×10^{-5}	1986 12 04	1989 06 09	4.998×10^{-4}	1.859×10^{-4}
		1994 11 22	1995 12 19	1.054×10^{-3}	2.012×10^{-5}	1990 08 02	1990 11 09	-1.948×10^{-3}	1.538×10^{-4}
		2061 Days				1580 Days			

The table presents the dates, mean return, and unconditional variance of the return for each of the regimes identified in the raw return series. The base regime covers the entire sample period from January 2, 1952 through December 29, 1995 (11,092 observations), with the returns for high and low volatility regimes excluded from the calculation of the sample statistics. Each row in the sections of the table labeled "Low Volatility" and "High Volatility" represent a single period identified as either a regime of unusually low or unusually high unconditional volatility respectively. The number of days identified at the end of each column identifies the total number of days identified with that particular type of regime in the sample period.

Table 6.8: Dates of regimes 1915–1929

Base		Low Volatility				High Volatility			
$\bar{r}$	Var	From	To	$\bar{r}$	Var	From	To	$\bar{r}$	Var
4.670×10^{-4}	7.087×10^{-5}	1918 06 05	1918 09 10	5.355×10^{-3}	2.289×10^{-4}	1917 10 26	1918 02 18	5.083×10^{-3}	2.034×10^{-3}
3722 Days		1922 01 12	1922 03 31	1.711×10^{-2}	2.365×10^{-4}	91 Days			
		1922 12 19	1923 03 27	8.015×10^{-3}	2.272×10^{-4}				
		1923 11 20	1924 02 16	6.404×10^{-3}	3.218×10^{-4}				
		1924 06 07	1924 09 06	1.536×10^{-2}	2.361×10^{-4}				
		1925 06 09	1925 09 03	9.420×10^{-3}	2.392×10^{-4}				
		1926 10 28	1927 01 26	3.088×10^{-3}	2.238×10^{-4}				
		1927 03 17	1927 06 04	1.010×10^{-2}	2.494×10^{-4}				
		593 Days							

The table presents the dates, mean return, and unconditional variance of the return for each of the regimes identified in the raw return series. The base regime covers the entire sample period from January 2, 1915 through September 30, 1929 (4,406 observations), with the returns for high and low volatility regimes excluded from the calculation of the sample statistics. Each row in the sections of the table labeled "Low Volatility" and "High Volatility" represent a single period identified as either a regime of unusually low or unusually high unconditional volatility respectively. The number of days identified at the end of each column identifies the total number of days identified with that particular type of regime in the sample period.

Table 6.9: Dates of regimes 1929–1951

Base		Low Volatility				High Volatility			
$\bar{r}$	Var	From	To	$\bar{r}$	Var	From	To	$\bar{r}$	Var
-5.496×10^{-4}	8.093×10^{-3}	1929 12 04	1930 06 17	-1.787×10^{-3}	1.946×10^{-3}	1929 10 01	1951 12 31	*none present*	*none present*
2152 Days		1930 06 21	1930 10 02	-3.481×10^{-3}	2.341×10^{-3}	0 Days			
		1930 12 15	1931 05 01	2.148×10^{-3}	2.219×10^{-3}				
		1931 06 25	1931 09 19	-3.715×10^{-2}	2.583×10^{-3}				
		1933 10 31	1937 10 19	2.869×10^{-3}	1.671×10^{-3}				
		1938 06 30	1938 09 14	-3.736×10^{-3}	2.888×10^{-3}				
		1938 09 23	1939 01 24	2.951×10^{-3}	2.151×10^{-3}				
		1939 04 08	1950 06 27	3.544×10^{-3}	8.370×10^{-4}				
		1950 06 29	1951 12 31	8.794×10^{-3}	6.022×10^{-4}				
		5490 Days							

The table presents the dates, mean return, and unconditional variance of the return for each of the regimes identified in the raw return series. The base regime covers the entire sample period from October 10, 1929 through December 31, 1951 (6,536 observations), with the returns for high and low volatility regimes excluded from the calculation of the sample statistics. Each row in the sections of the table labeled "Low Volatility" and "High Volatility" represent a single period identified as either a regime of unusually low or unusually high unconditional volatility respectively. The number of days identified at the end of each column identifies the total number of days identified with that particular type of regime in the sample period.

Table 6.10: Dates of regimes 1952–1973

Base		Low Volatility				High Volatility			
$\bar{r}$	Var	From	To	$\bar{r}$	Var	From	To	$\bar{r}$	Var
5.214×10^{-4}	4.210×10^{-5}	1952 05 03	1952 10 16	3.825×10^{-4}	1.618×10^{-5}	1955 08 22	1955 12 14	1.198×10^{-3}	1.413×10^{-04}
3854 Days		1957 03 04	1957 05 31	1.352×10^{-3}	1.447×10^{-5}	1957 08 06	1957 12 12	-1.533×10^{-3}	1.301×10^{-04}
		1963 02 26	1963 11 22	5.445×10^{-4}	1.633×10^{-5}	1962 03 23	1962 10 11	-1.373×10^{-3}	1.495×10^{-04}
		1963 11 27	1966 05 10	4.411×10^{-4}	1.572×10^{-5}	1970 02 16	1970 08 06	-1.169×10^{-3}	1.436×10^{-04}
		1967 06 07	1967 11 02	4.834×10^{-4}	1.561×10^{-5}	520 Days			
		1968 07 26	1968 12 31	9.789×10^{-4}	1.577×10^{-5}				
		1185 Days							

The table presents the dates, mean return, and unconditional variance of the return for each of the regimes identified in the raw return series. The base regime covers the entire sample period from January 2, 1952 through December 31, 1973 (5,531 observations), with the returns for high and low volatility regimes excluded from the calculation of the sample statistics. Each row in the sections of the table labeled "Low Volatility" and "High Volatility" represent a single period identified as either a regime of unusually low or unusually high unconditional volatility respectively. The number of days identified at the end of each column identifies the total number of days identified with that particular type of regime in the sample period.

Table 6.11: Dates of regimes 1974–1995

Base		Low Volatility				High Volatility			
$\bar{r}$	Var	From	To	$\bar{r}$	Var	From	To	$\bar{r}$	Var
4.386×10^{-4}	6.043×10^{-5}	1992 10 05	1993 02 17	7.932×10^{-4}	2.578×10^{-5}	1974 06 12	1975 04 25	-1.164×10^{-04}	1.964×10^{-04}
4069 Days		1993 03 09	1994 02 07	2.814×10^{-4}	2.157×10^{-5}	1982 07 30	1982 11 12	4.062×10^{-03}	1.874×10^{-04}
		1994 05 12	1994 08 29	9.613×10^{-4}	2.093×10^{-5}	1986 12 04	1988 10 31	2.880×10^{-04}	2.332×10^{-04}
		1994 11 22	1995 12 29	1.085×10^{-3}	1.999×10^{-5}	779 Days			
		683 Days							

The table presents the dates, mean return, and unconditional variance of the return for each of the regimes identified in the raw return series. The base regime covers the entire sample period from January 2, 1974 through December 29, 1995 (5,559 observations), with the returns for high and low volatility regimes excluded from the calculation of the sample statistics. Each row in the sections of the table labeled "Low Volatility" and "High Volatility" represent a single period identified as either a regime of unusually low or unusually high unconditional volatility respectively. The number of days identified at the end of each column identifies the total number of days identified with that particular type of regime in the sample period.

Table 6.12: Estimates of model parameters (sub-sample, no regimes)

Model	γ_0	γ_2	γ_3	γ_4	λ	ν	b	ω	α	β
1915–1995	$6.297^{[-4]}$ $(5.155^{[-5]})^*$	0.145 (0.008)*	−0.023 (0.008)*	−0.002 $(1.290^{[-4]})^*$	1.167 (0.276)*	1.611 (0.112)*	0.532 (0.050)*	$7.015^{[-5]}$ $(9.282^{[-5]})$	0.069 (0.016)*	0.910 (0.020)*
1915–1951	$6.434^{[-4]}$ $(8.734^{[-5]})^*$	0.078 (0.011)*	−0.017 (0.011)	−0.002 $(2.341^{[-4]})^*$	2.850 (0.744)*	2.007 (0.201)*	0.651 (0.063)*	$4.100^{[-8]}$ $(1.472^{[-7]})$	0.150 (0.041)*	0.785 (0.050)*
1952–1995	$6.227^{[-4]}$ $(6.343^{[-5]})^*$	0.214 (0.012)*	−0.039 (0.013)*	−0.002 $(1.589^{[-4]})^*$	0.247 (0.504)	1.192 (0.166)*	0.512 (0.128)*	0.006 (0.016)	0.019 (0.039)	0.961 (0.038)*
1915–1929	$6.923^{[-4]}$ $(1.247^{[-4]})^*$	0.088 (0.019)*	0.008 (0.017)	−0.001 $(3.207^{[-4]})^*$	5.531 (0.063)*	3.030 (0.312)*	1.376 (0.183)*	$1.568^{[-13]}$ $(4.156^{[-14]})^*$	0.107 (0.044)*	0.388 (0.071)*
1929–1951	$5.567^{[-4]}$ $(1.253^{[-4]})^*$	0.087 (0.015)*	−0.033 (0.014)*	−0.002 $(3.330^{[-4]})^*$	1.648 (0.807)*	1.757 (0.214)*	0.585 (0.078)*	$1.120^{[-5]}$ $(4.186^{[-5]})$	0.093 (0.044)*	0.877 (0.051)*
1952–1973	$7.160^{[-4]}$ $(8.611^{[-5]})^*$	0.273 (0.015)*	−0.073 (0.015)*	−0.002 $(2.055^{[-4]})^*$	$7.863^{[-4]}$ (0.137)	0.970 (0.150)*	0.515 (0.063)*	0.046 (0.003)*	$1.095^{[-4]}$ (0.019)	0.953 (0.048)*
1974–1995	$4.586^{[-4]}$ $(1.057^{[-4]})^*$	0.161 (0.014)*	−0.012 (0.014)	$-4.686^{[-4]}$ $(2.461^{[-4]})$	2.826 (2.107)	2.027 (0.735)*	0.503 (0.093)*	$1.718^{[-8]}$ $(1.738^{[-7]})$	0.091 (0.052)	0.878 (0.091)*

Entries of the form $x^{[-y]}$ represent scientific notation; that is $2.5^{[-2]}$ represents the value 2.5×10^{-2}.

The table presents the subsample parameter estimates for the conventional GARCH model. Numbers in parentheses under the parameter estimates represent asymptotic standard errors. Parameters which are significantly different from zero at the 5% level are indicated with an asterisk.

Table 6.13: Estimates of model parameters (sub-sample)

Model	γ_0	γ_2	γ_3	γ_4	λ	ν	b	ω	α	β
1915–1995	$6.780^{[-4]}$	0.147	−0.025	−0.001	3.149	2.295	0.828	$1.769^{[-8]}$	0.098	0.725
	$(5.019^{[-5]})^*$	$(0.008)^*$	$(0.007)^*$	$(1.250^{[-4]})^*$	$(1.356)^*$	$(0.365)^*$	$(0.065)^*$	$(1.179^{[-7]})$	$(0.028)^*$	$(0.078)^*$
					−1.314	−0.861	0	$4.220^{[-6]}$	0.071	0
					(0.766)	$(0.338)^*$		$(1.925^{[-5]})$	(0.038)	
					0.198	0	0	0	0.045	0
					(0.148)				$(0.010)^*$	
1915–1951	$7.022^{[-4]}$	0.084	−0.017	−0.001	4.495	3.069	1.637	$4.891^{[-11]}$	0.046	0.477
	$(7.902^{[-5]})^*$	$(0.011)^*$	(0.010)	$(8.678^{[-5]})^*$	$(1.105)^*$	$(0.308)^*$	$(0.047)^*$	$(2.743^{[-10]})$	$(0.017)^*$	$(0.065)^*$
					−3.394	−2.524	0	$7.352^{[-4]}$	0.166	0
					$(1.091)^*$	$(0.349)^*$		$(3.293^{[-4]})^*$	$(0.042)^*$	
					−3.818	−1.430	−0.866	0.002	0	0.433
					$(1.071)^*$	$(0.402)^*$	$(0.103)^*$	(0.002)		$(0.058)^*$
1952–1995	$6.760^{[-4]}$	0.212	−0.036	−0.001	5.212	2.369	0.799	$2.010^{[-13]}$	0.117	0.743
	$(6.387^{[-5]})^*$	$(0.011)^*$	$(0.010)^*$	$(1.474^{[-4]})^*$	$(1.716)^*$	$(0.369)^*$	$(0.095)^*$	$(1.741^{[-12]})$	$(0.032)^*$	$(0.067)^*$
					−3.998	−0.859	0	$1.464^{[-4]}$	0	0
					$(1.455)^*$	$(0.336)^*$		$(4.371^{[-4]})$		
					−2.053	0	0	$3.195^{[-8]}$	0	0
					(1.271)			$(1.603^{[-7]})$		

Entries of the form $x^{[-y]}$ represent scientific notation; that is $2.5^{[-2]}$ represents the value 2.5×10^{-2}.

Table 6.13: continued.

Model	γ_0	γ_2	γ_3	γ_4	λ	ν	b	ω	α	β
1915–1929	$7.304^{[-4]}$	0.088	0.002	−0.001	6.993	8.897	6.225	$1.923^{[-16]}$	$4.404^{[-8]}$	0.149
	$(1.206^{[-4]})^*$	(0.018)*	(0.017)	$(3.129^{[-4]})^*$	(4.155)	(6.918)	(4.195)	$(4.007^{[-15]})$	$(7.958^{[-7]})$	(0.162)
					−0.505	0	0	0	0	0
					(0.291)					
					0	0.138	0	0	0	0
						(0.068)*				
1929–1951	$5.607^{[-4]}$	0.087	−0.036	−0.002	0.345	1.535	0.626	0.010	0.021	0.942
	$(1.238^{[-4]})^*$	(0.014)*	(0.014)*	$(3.302^{[-4]})^*$	(0.672)	(0.205)*	(0.086)*	(0.024)	(0.039)	(0.046)*
					0	0	0	−0.003	0	0
								(0.002)		
					0	0	0	0	0	0
1952–1973	$7.348^{[-4]}$	0.278	−0.066	−0.002	$7.512^{[-4]}$	1.523	0.811	0.119	$5.660^{[-5]}$	0.881
	$(7.519^{[-5]})^*$	(0.015)*	(0.014)*	$(1.734^{[-4]})^*$	$(2.644^{[-4]})^*$	(0.170)*	(0.125)*	(0.023)*	$(2.652^{[-5]})^*$	(0.023)*
					0	0	0	0	0	$-4.465^{[-5]}$
										$(1.338^{[-5]})^*$
					0	0	0	0	$-2.020^{[-5]}$	$7.918^{[-5]}$
									$(6.838^{[-6]})^*$	$(1.117^{[-7]})^*$
1974–1995	$4.955^{[-4]}$	0.161	−0.016	$-4.505^{[-4]}$	3.123	1.989	0.641	$1.130^{[-8]}$	0.065	0.863
	$(1.016^{[-4]})^*$	(0.015)*	(0.014)	$(2.411^{[-4]})$	(1.221)*	(0.365)*	(0.061)*	$(7.542^{[-8]})$	(0.034)	(0.001)*
					0	0	2.778	0	0	−0.863
							(0.001)*			$(1.597^{[-6]})^*$
					0	0	0	0	0.161	−0.125
									$(1.634^{[-5]})^*$	$(6.842^{[-7]})^*$

Entries of the form $x^{[-y]}$ represent scientific notation; that is $2.5^{[-2]}$ represents the value 2.5×10^{-2}.

The table presents the subsample parameter estimates for the Regime Dependent GARCH model. Numbers in parentheses under the parameter estimates represent asymptotic standard errors. Parameters which are significantly different from zero at the 5% level are indicated with an asterisk. The first pair of rows in each section present the parameter estimates and standard errors for the base regime, and the remaining four rows present the parameter estimates and standard errors for the regimes of unusually low and high volatility respectively. Parameters which are excluded by the stepwise estimation procedure are represented by a value of zero, with no corresponding standard error presented.

Chapter 7

Out of Sample Forecasting

During development of the Regime Dependent GARCH model, the objective was improving within-sample model performance, as measured by the production of higher likelihood values. The selection of the final model, made on the basis of likelihood ratio tests, provided a model superior to the others investigated at explaining the within-sample volatility of the stock return index. While a model that performs well within-sample is useful for describing and perhaps aiding in understanding the effects that are present within a particular dataset, the ability of the model to predict future events can be of great benefit in decision making. The purpose of this chapter is to determine how well the RD-GARCH model of conditional volatility performs in predicting future (out-of-sample) variance. The chapter begins with a brief note on the notational conventions used. The second section discusses the methodology used in constructing the forecasts. The third section outlines the series that will be used as proxies for the *ex post* realized volatility, and the fourth section identifies several models which will be used to provide alternative forecasts of future variance for comparative purposes. Finally, the fifth section outlines the tests that will be performed on the various forecasts and the sixth section discusses the results of the tests.

7.1 A Note on Notation

In this chapter several types of forecasts will be constructed, discussed, and analyzed. For clarity, the following notation will be adopted: *Ex-post* realized variance at time t will be represented by ϑ_t, with daily variance represented by ϑ_t^d and monthly variance represented by ϑ_t^m. A variance forecast made at time t will be represented by $\widehat{\vartheta}_t$, with the subscript identifying also the type of forecast and the number of periods ahead for which the forecast is made. The superscript will indicate whether daily or monthly variances have been forecast. Thus, $\widehat{\vartheta}^d_{u,t+1}$ will represent a one-step ahead forecast of daily variance using the unconditional variance model. Finally, the forecast error obtained under a particular model will be denoted by: $e_t = \vartheta_t - \widehat{\vartheta}_t$ with the superscript again identifying whether daily or monthly forecasts have been made, and the subscript identifying the type of forecast and the number of periods ahead for which the forecast has been made. Thus, $e^d_{u,t+1} = \vartheta^d_{u,t+1} - \widehat{\vartheta}^d_{u,t+1}$ represents the one-period ahead daily variance forecast error using the unconditional variance model.

7.2 Methodology

A general approach to conducting tests of forecast performance proceeds as follows: The model is first calibrated by estimating the parameters using a portion of the available data. The parameter estimates obtained during the calibration phase are then used to predict future values for the series. The forecast values are finally compared to the actual values from the portion of the series that was not used for calibration.

The previous chapters developed the Regime Dependent GARCH model using data covering the period January 2, 1915 to December 29, 1995. Although the daily CRSP index series that comprises the latter portion of the dataset is currently available to December 31, 1997, the final two years of the dataset were retained as a hold-out sample to assess the out-of-sample forecasting performance of the RD-GARCH model.

A common approach to calibration of forecasting models involves the use of a rolling calibration period (See, for example, Akgiray, 1989; Lamoureaux and Lastrapes, 1993; Brooks, 1998). A fixed number of observations, N_r, is maintained as the calibration period. The model is estimated over the calibration period, and the n-step ahead forecast[1] is made using the parameters estimates obtained during calibration. The realized values from the dataset which overlap those forecast are then added to the end of the calibration sample, and the same number of observations are removed from the beginning of the period so that the total length of the calibration period remains constant at N_r. The model is re-estimated for the new calibration period, new forecasts are made, and the process repeats.

An alternative approach uses an updating calibration period (Lamoureaux and Lastrapes, 1993, use both rolling and updating periods to calibrate models). In this approach, a fixed starting date for the calibration period is maintained. The model is estimated for the calibration period, and the forecast is made. The realized values from the dataset which overlap those forecast are then added to the end of the calibration sample, but no data are removed from the beginning of the period, so the total length of the calibration period grows with each successive forecast. Again, the model is re-estimated for the new calibration period, new forecasts are made, and the process repeats.

Care must be taken in selecting the calibration period. If it is too short, the model parameter values obtained from the calibration may merely reflect recent transient phenomena. If the period selected is too long, it may span structural changes in the underlying series which will cause the resulting parameter values not to reflect accurately the current state of the world at the time the forecast is made. From the results of the sub-period investigations in Chapter 6, two potential updating calibration periods are suggested: the first period begins January 2, 1952, and the second begins January 2, 1973. A third alternative involving a shorter five-year rotating calibration period is also investigated.

---

[1] Where n is frequently one.

Two forecast periods are also considered. Because the dataset used initially to develop and now to calibrate the model comprises daily returns, the raw forecasts produced are of daily volatility. Practical applications of volatility models might well require estimates of volatility over periods longer than a single day, however. For these reasons, forecasts of both daily and monthly volatility are constructed and evaluated.

One additional consideration involves the frequency of calibration. As outlined above, the usual procedure is to calibrate the model, generate a one-step ahead forecast (plus possibly additional n-step ahead forecasts), adjust the calibration period, and repeat the entire process. Unfortunately, estimation of each RD-GARCH model can take several days of computer time, so rather than re-estimating the model each time a forecast is made, calibration is performed quarterly to estimate new model parameters. The one-period ahead daily variance forecasts for $t+1$ are made using the parameter estimates from the most recent calibration and the data from the stock return series up to time t. Similarly, the one-period ahead monthly variance forecasts for $t+1$ are made using the parameter estimates from the most recent calibration and the data from the stock return series up to time t to generate forecasts of the one-period ahead through N_m-period ahead daily variances (where N_m is the number of trading days in the month). The daily variances are summed to provide an estimate of the monthly variance. Although not employed in his work, Akgiray (1989) suggests such a procedure.

7.3 The Volatility of Returns

To test the effectiveness of a model in predicting out-of sample variances, it is necessary to provide a proxy for the true volatility of the return series. As the GARCH models have been calibrated with daily data, they will generate forecasts of daily conditional volatility. Where transactions level data are available, or where there is information regarding the high and low stock price, there are several alternative methods of computing estimates of the daily variance of the return series. As the long-term dataset used in this study provides only the closing level of the index, the approach of Day and Lewis (1992) and Brooks (1998) using the squared daily return as a proxy for the daily variance is employed. The daily variance on the stock index is given by:

$$\vartheta_t^d = r_t^2 \tag{7.1}$$

The monthly variance on the stock index can be constructed from the daily data in a number of ways. The simplest approach is to compute the monthly return, r_{tm}, given by:

$$r_{tm} = \sum_{i=1}^{N_m} r_{t+i-1} \tag{7.2}$$

The monthly variance is then computed in a manner identical to the method used to compute the daily variance:

$$\vartheta_t^m = r_{tm}^2 \tag{7.3}$$

As daily returns are available, an alternative approach is available which makes use of the well known property:

$$\operatorname{var}(x_1 + x_2) = \operatorname{var}(x_1) + \operatorname{var}(x_2) + 2 \cdot \operatorname{cov}(x_1, x_2) \tag{7.4}$$

which generalizes to:

$$\vartheta_t^m = \sum_{i=1}^{N_m} r_{t+i-1}^2 + 2 \sum_{i=1}^{N_m-1} \sum_{j=i+1}^{N_m} r_{t+i-1} r_{t+j-1} \tag{7.5}$$

Some researchers (eg. Day and Lewis, 1992; Brooks, 1998) suggest that the cross-product terms may be dropped, making an implicit assumption that the returns exhibit no autocorrelation. Others (eg. French et al., 1987; Akgiray, 1989) suggest that since stock portfolio returns tend to be autocorrelated, particularly at lag one, it is appropriate to include the term associated with adjacent returns:

$$\vartheta_t^m = \sum_{i=1}^{N_m} r_{t+i-1}^2 + 2 \sum_{i=1}^{N_m-1} r_{t+i-1} r_{t+i} \tag{7.6}$$

For the purposes of forecast evaluation of monthly variances, results for using the monthly variance from Equation 7.6 as a proxy for the *ex post* realized "true" monthly variance are presented.

7.4 Forecasting Models

While the aim of this chapter is to determine how well the RD-GARCH model performs at out-of-sample variance prediction, it is not possible to make such a determination in isolation. Instead, forecasts are generated using an unconditional variance model, several moving average models, a random walk model, a conventional GARCH model, the Extended GARCH model developed in Chapter 4, and the RD-GARCH model. Tests are then performed to determine which forecasting model is preferred.

Squared returns are forecast as a proxy for return variance in the unconditional variance, moving average, and random walk forecast models. The use of squared returns makes the simplifying assumption that $E(r_t) = 0$. As the proxy for *ex post* realized return variance outlined in Section 7.3 makes the same assumption, this simplification seems reasonable. The conditional variance of the residual series from the conditional mean equation is used as a proxy for return variance in the GARCH-based models.[2]

7.4.1 Unconditional Variance

The simplest forecast of the future variance is the unconditional variance of the entire realized return series. If variance were indeed constant, this would represent the optimal prediction of future variance. The daily one-step ahead forecast using the unconditional variance model is given by:

$$\widehat{\vartheta}^{d}_{u,t+1} = \frac{1}{N-1}\sum_{t=1}^{N} r_t^2 \tag{7.7}$$

Where $t = 1$ corresponds to the beginning of the calibration period, and N corresponds to the number of daily observations in the calibration period. As the unconditional variance is not time dependent, the n-step ahead forecast is simply:

$$\widehat{\vartheta}^{d}_{u,t+n} = \widehat{\vartheta}^{d}_{u,t+1} = \frac{1}{N-1}\sum_{t=1}^{N} r_t^2 \tag{7.8}$$

The one-period ahead forecast of monthly variance is given by:

$$\widehat{\vartheta}^{m}_{u,t+1} = N_m \cdot \widehat{\vartheta}^{d}_{u,t+1} \tag{7.9}$$

7.4.2 Moving Average Variance

An alternative to the unconditional variance model, which does not rely on constant variance, is a moving average of prior squared returns. For comparative purposes, three moving average window sizes of five days, 20 days, and 60 days (representing a trading week, trading month, and trading quarter respectively) are estimated. The daily one-step ahead forecast using the moving average variance model is given by:

---

[2] The residual series is defined as $u_t = r_t - E(r_t|I_{t-1})$ where $E(r_t|I_{t-1})$ is the conditional mean of the return series from Equation 4.11. The conditional variance of u_t is $\mathrm{var}(u_t|I_{t-1}) = E(r_t^2|I_{t-1}) - E(r_t|I_{t-1})^2 = h_t^2$, while the proxy for the variance of the return series, r_t, outlined in Section 7.3 is $\vartheta_t = r_t^2$. Thus, the use of h_t^2 as a proxy for the conditional variance of r_t introduces a bias into the forecast. The average daily return is empirically very small, however, so this bias will also be small. As this is the approach which is commonly used, and in the interests of simplicity, the approach is adopted here in spite of the bias.

$$\widehat{\vartheta}^d_{mL,t+1} = \frac{1}{L-1}\sum_{i=1}^{L} r^2_{N-i+1} \tag{7.10}$$

Where L corresponds to the window size, and $\widehat{\vartheta}^d_{m20,t+1}$ represents the daily one-step ahead forecast using a 20 day moving average.[3] As the moving average variance uses only observed returns, the n-step ahead forecast is again simply:

$$\widehat{\vartheta}^d_{mL,t+n} = \widehat{\vartheta}^d_{mL,t+1} = \frac{1}{L-1}\sum_{i=1}^{L} r^2_{N-i+1} \tag{7.11}$$

and the forecast monthly variance is given by:

$$\widehat{\vartheta}^m_{mL,t+1} = N_m \cdot \widehat{\vartheta}^d_{mL,t+1} \tag{7.12}$$

7.4.3 Random Walk

A random walk model for return variance is given by:

$$\vartheta^d_t = r^2_{t-1} + \varepsilon_t \tag{7.13}$$

The daily one-step ahead forecast using the random walk model is given by:

$$\widehat{\vartheta}^d_{r,t+1} = r^2_t \tag{7.14}$$

The random walk model is another special case of the moving average variance model with $\widehat{\vartheta}^d_{r,t+1} \equiv \widehat{\vartheta}^d_{m1,t+1}$; tomorrow's variance is forecast as being the same as today's. Again, the model uses only observed returns, so the n-step ahead forecast is simply:

$$\widehat{\vartheta}^d_{r,t+n} = \widehat{\vartheta}^d_{r,t+1} = r^2_t \tag{7.15}$$

and the forecast monthly variance is given by:

$$\widehat{\vartheta}^m_{r,t+1} = N_m \cdot \widehat{\vartheta}^d_{r,t+1} \tag{7.16}$$

---

[3] Of course, the unconditional variance model is a special case of the moving average variance model with $\widehat{\vartheta}^d_{u,t+1} \equiv \widehat{\vartheta}^d_{mN,t+1}$.

7.4.4 Conventional GARCH

The conventional GARCH(1,1) model provides forecasts of the conditional variance, $\hat{h}^2_{t+i}$, according to the relationship given previously in Equation 1.7:

$$h_t^2 = \omega + \alpha u_{t-1}^2 + \beta h_{t-1}^2 \tag{7.17}$$

where the model parameters are as outlined in Section 1.2.2, and h_t^2, the conditional variance of u_t, is used as a proxy for the conditional variance of r_t. Making use of the relationship from Equation 1.4 which defines the standardized residual, $v_t = u_t/h_t \sim iid(0,1)$, and rearranging yields:

$$\begin{aligned} h_t^2 &= \omega + \alpha h_{t-1}^2 v_{t-1}^2 + \beta h_{t-1}^2 && (7.18a) \\ &= \omega + \alpha h_{t-1}^2 \left(v_{t-1}^2 - 1\right) + (\alpha + \beta)\, h_{t-1}^2 && (7.18b) \end{aligned}$$

The daily one-step ahead forecast using the GARCH(1,1) model is given by:

$$\widehat{\vartheta}^d_{G,t+1} = \hat{h}^2_{t+1} = \omega + \alpha h_t^2 \left(v_t^2 - 1\right) + (\alpha + \beta)\, h_t^2 \tag{7.19}$$

where h_t^2 and v_t^2 are both known at time t. The conditional variance is not time invariant, so the n-step ahead forecast is given by:

$$\widehat{\vartheta}^d_{G,t+n} = \hat{h}^2_{t+n} = \omega + \alpha h^2_{t+n-1} \left(v^2_{t+n-1} - 1\right) + (\alpha + \beta)\, h^2_{t+n-1} \tag{7.20}$$

If the forecasting is performed in a sequential manner (ie. if the forecast of $\hat{h}^2_{t+n}$ is made following the forecasts of $\hat{h}^2_{t+n-i}$, $i = 1 \ldots n$), the forecast value of $\hat{h}^2_{t+n-1}$ will be known at time $t + n$. The value of the unexpected shock, v_{t+n}, $n > 0$, is not knowable prior to $t + n$, and is replaced by its unconditional expectation:

$$v_t \sim iid(0,1), \qquad E(v_t^2) = 1 \tag{7.21}$$

Equation 7.20 reduces to:

$$\widehat{\vartheta}^d_{G,t+n} = \omega + (\alpha + \beta)\, \widehat{\vartheta}^d_{G,t+n-1}, \qquad n > 1 \tag{7.22}$$

The one-period ahead forecast of monthly variance is then computed as the sum of N_m individual n-step ahead forecasts:

$$\widehat{\vartheta}^m_{G,t+1} = \sum_{n=1}^{N_m} \widehat{\vartheta}^d_{G,t+n} \tag{7.23}$$

where the individual $\widehat{\vartheta}^d_{G,t+n}$ terms have been computed by repeated application of Equation 7.22.

7.4.5 Extended GARCH

The Extended GARCH model developed in Chapter 4 provides forecasts of conditional volatility, $\hat{h}_{t+i}$, in a manner analogous to the GARCH model, although the specification is somewhat more involved. The conditional volatility relationship was given previously in Equation 4.12:

$$h_t = \left(\omega + \alpha h_{t-1}^{\lambda}|v_{t-1} - b|^{\nu} + \beta h_{t-1}^{\lambda}\right)^{\frac{1}{\lambda}} \tag{7.24}$$

where the model parameters are as outlined in Sections 1.2.4 and 1.3, and h_t^2 is again the conditional variance of the return series. The daily one-step ahead forecast of the conditional variance using the Extended GARCH model is given by:

$$\widehat{\vartheta}^{d}_{E,t+1} = \hat{h}^2_{t+1} = \left(\omega + \alpha h_t^{\lambda}|v_t - b|^{\nu} + \beta h_t^{\lambda}\right)^{\frac{2}{\lambda}} \tag{7.25}$$

where both h_t and v_t are known at time t. Again, the conditional variance is time dependent, so the n-step ahead forecast is given by:

$$\widehat{\vartheta}^{d}_{E,t+n} = \hat{h}^2_{t+n} = \left(\omega + \alpha \hat{h}_{t+n-1}^{\lambda}|v_{t+n-1} - b|^{\nu} + \beta \hat{h}_{t+n-1}^{\lambda}\right)^{\frac{2}{\lambda}} \tag{7.26}$$

As was the case with the conventional GARCH model, if the forecasting is performed in a sequential manner the value of $\hat{h}_{t+n-1}$ will be known at time $t+n$ but the value of the unexpected shock, v_{t+n-1} will be unknown. Unfortunately, the term in Equation 7.26 analogous to the $(v_{t+n-1}^2 - 1)$ term in Equation 7.20 is $|v_{t+n-1} - b|^{\nu}$, for which a usable general expression for the unconditional expectation does not exist.[4] In order to provide forecasts for the conditional variance, the mean of $|v_{t+n-1} - b|^{\nu}$, denoted $\bar{f}$, is computed during the calibration period and used as a proxy for its unconditional expectation during the forecast period. Equation 7.26 can then be rewritten as:

$$\begin{aligned}\widehat{\vartheta}^{d}_{E,t+n} &= \left(\omega + \alpha \hat{h}_{t+n-1}^{\lambda}\bar{f} + \beta \hat{h}_{t+n-1}^{\lambda}\right)^{\frac{2}{\lambda}} \\ &= \left(\omega + \left(\alpha \bar{f} + \beta\right)\hat{h}_{t+n-1}^{\lambda}\right)^{\frac{2}{\lambda}} \end{aligned} \tag{7.27}$$

As is the case for the conventional GARCH model, the one-period ahead forecast of monthly variance is then computed as the sum of N_m individual n-step ahead forecasts:

---

[4] Under the assumption of normality, the unconditional expectation of $|v_{t+n-1} - b|^{\nu}$ was investigated using the Maple V (Waterloo Maple Inc., 1998) symbolic mathematics system. A simple functional form for the expectation does exist when ν takes on values which are even whole numbers, (for $\nu = 2$, for example, $E\left[|v_{t+n-1} - b|^{\nu}\right] = b^2 + 1$). Odd whole numbers provide a functional form which is defined in terms of $\operatorname{erf}\left(b/\sqrt{2}\right)$, which does provide numeric solutions when the value of b is specified. Values of ν which are not whole numbers yield a functional form which provides numeric solutions for specified values of b, however these solutions $\in \mathbb{C}$.

$$\widehat{\vartheta}^{m}_{E,t+1} = \sum_{n=1}^{N_m} \widehat{\vartheta}^{d}_{E,t+n} \tag{7.28}$$

with the individual $\widehat{\vartheta}^{d}_{E,t+n}$ terms computed from repeated application of Equation 7.27.

7.4.6 RD-GARCH

The Regime Dependent GARCH model obtains from the Extended GARCH specification, and provides additional flexibility by allowing the parameters to vary across regimes of unusually high or low volatility. The model was developed in Chapter 4 and, like the Extended GARCH model, provides estimates of the conditional volatility, h_t, although the specification is somewhat more involved. The conditional volatility relationship given previously in Equation 4.17 was:

$$h_t = \left(\boldsymbol{\omega} + \boldsymbol{\alpha} h_{t-1}^{\boldsymbol{\lambda}} |v_{t-1} - \boldsymbol{b}|^{\boldsymbol{\nu}} + \boldsymbol{\beta} h_{t-1}^{\boldsymbol{\lambda}}\right)^{\frac{1}{\boldsymbol{\lambda}}} \tag{7.29}$$

where each of the parameters indicated in bold is conditioned by two $(0, 1)$ dummy variables, D_L and D_H, which are zero in periods of normal volatility and one in periods of low or high volatility respectively. That is, each parameter has been allowed to vary from the base regime for both low and high volatility regimes. For example, λ in Equation 7.24 is replaced by:

$$\boldsymbol{\lambda} = \lambda_B + \lambda_L D_L + \lambda_H D_H \tag{7.30}$$

in Equation 7.29. Forecasting from this model follows essentially the same process as forecasting from the simpler Extended GARCH specification outlined above in Equations 7.25 and 7.27, except that during the forecasting process an assessment must be made regarding the appropriate current volatility regime to determine the values for the parameters of the model. As the regime detection and isolation process outlined in Section 4.2.3 used a moving window of 60 days (representing a trading quarter) to isolate periods of unusually high or low volatility, the forecast computes the variance of the prior 30 trading days of *ex post* realized returns and constructs a variance ratio, $\hat{R}_d$, comparing the variance of these 30 returns to the variance of the base regime from the calibration period. If the variance ratio exceeds the critical value of 3.0 used in the original detection algorithm, the parameter estimates for high volatility regimes are used to construct the forecast. Alternatively, if $\hat{R}_d^{-1}$ exceeds the critical value, the parameter estimates for low volatility regimes are used to construct the forecast. If neither $\hat{R}_d$ or $\hat{R}_d^{-1}$ exceeds the critical value, the parameter estimates for the base regime are used to construct the forecast. The daily one-step ahead forecast of the conditional variance using the RD-GARCH model is given by:

$$\widehat{\vartheta}^{d}_{D,t+1} = \hat{h}^2_{t+1} = \left(\boldsymbol{\omega} + \boldsymbol{\alpha} h_t^{\boldsymbol{\lambda}} |v_t - \boldsymbol{b}|^{\boldsymbol{\nu}} + \boldsymbol{\beta} h_t^{\boldsymbol{\lambda}}\right)^{\frac{2}{\boldsymbol{\lambda}}} \tag{7.31}$$

where both h_t and v_t are known at time t and an assessment of the current regime has been made based on the variance of the previous 30 known returns. Again, the conditional variance is time dependent, so the n-step ahead forecast is given by:

$$\widehat{\vartheta}^d_{D,t+n} = \hat{h}^2_{t+n} = \left(\omega + \alpha \hat{h}^{\lambda}_{t+n-1}|v_{t+n-1} - b|^{\nu} + \beta \hat{h}^{\lambda}_{t+n-1}\right)^{\frac{2}{\lambda}} \tag{7.32}$$

As was the case with the Extended GARCH model, if forecasting is performed in a sequential manner, the value of $\hat{h}_{t+n-1}$ will be known at time $t+n$ but the value of the unexpected shock, v_{t+n-1} will be unknown. The problem does increase in complexity somewhat due to the regime dependent nature of the parameters of the model. The solution adopted follows the solution to the related problem in the extended GARCH model case. The mean of the $|v_{t+n-1} - b|^{\nu}$ expression is computed separately for each regime type during the calibration period, and the appropriate value is used as a proxy for the unconditional expectation during the forecast period. That is, three values are computed: $\bar{f}_B$ represents the mean of $|v_{t+n-1} - b|^{\nu}$ for those periods determined to be part of the base regime, $\bar{f}_L$ represents the mean for those periods determined to be part of a low volatility regime, and $\bar{f}_H$ represents the mean for those periods determined to be part of a high volatility regime. This approach does make the assumption that the regime type does not change during the forecast period, but as long as the number of periods ahead for which the forecasts are made is reasonably small the assumption should not be problematic. Equation 7.32 can then be rewritten as:

$$\widehat{\vartheta}^d_{D,t+n} = \left(\omega + \left(\alpha \bar{f} + \beta\right) \hat{h}^{\lambda}_{t+n-1}\right)^{\frac{2}{\lambda}} \tag{7.33}$$

where $\bar{f} = \bar{f}_L$, $\bar{f}_B$, or $\bar{f}_H$, if t is determined to belong to a low volatility regime, the base regime, or a high volatility regime respectively. As is the case for both the conventional and extended GARCH models, the one-period ahead forecast of monthly variance is computed as the sum of N_m individual n-step ahead forecasts:

$$\widehat{\vartheta}^m_{D,t+1} = \sum_{n=1}^{N_m} \widehat{\vartheta}^d_{D,t+n} \tag{7.34}$$

with the individual $\widehat{\vartheta}^d_{D,t+n}$ terms computed from repeated application of Equation 7.33.

7.5 Forecast Evaluation

Section 7.3 provided the method of constructing a daily and a monthly variance series to proxy for the *ex post* realized variance of the stock return index, and Section 7.4 outlined a number of alternatives for forecasting the volatility of the index. In this section, several alternative methods for evaluating the performance, both individually and comparatively, of the forecast methodologies are presented. These tests can be grouped into two general categories: loss function tests, and regression tests.

7.5.1 Loss Function Tests

In determining the effectiveness of a particular forecast, it is useful to think in terms of the difference between the *ex post* realized value, ϑ_t, and the forecast value, $\widehat{\vartheta}_{F,t}$. The difference between the realized value and the forecast value represents a loss in the sense that the forecast could be improved by reducing this difference. This loss can be summarized with a loss function which specifies the relative cost when the realized value is different from the forecast by a particular amount; the functional form can be specified in such a way as to penalize errors of one type more severely than errors of another type. Comparison of alternative forecast methodologies can then be made based upon the difference in total loss over the forecast period. The forecast which provides the smallest loss is the preferred forecast. There are two forms of loss function which are commonly used as criteria for assessing the relative performance of forecasts: linear loss functions and quadratic loss functions.

Mean Error

The mean error (ME) is not strictly a loss function based criterion for forecast evaluation in that it will provide an indication of bias inherent in a particular forecast, but will not provide information regarding the magnitude of the overall forecast error. While forecast bias is of interest and importance in comparing alternative forecast models, particularly when used in conjunction with the mean square error criterion outlined below, minimization of bias by itself is not a reasonable method of forecast evaluation. It is included in this section because of its computational similarity to the mean absolute error and mean square error criteria which follow. The ME over the forecast period is computed as:

$$\text{ME} = \frac{1}{N}\sum_{i=1}^{N}\left(\vartheta_i - \widehat{\vartheta}_{j,i}\right) = \frac{1}{N}\sum_{i=1}^{N} e_{j,i} = \bar{e}_j \tag{7.35}$$

where $\bar{e}_j$ represents the mean error of forecast model j. A positive ME indicates that on average the forecast underestimates the true variance. Similarly, a negative ME indicates that on average the forecast overestimates the true variance. Because the ME criterion represents the mean of the forecast error, the hypothesis of no bias can be investigated with traditional statistical inference techniques.

Mean Absolute Error

The mean absolute error (MAE) computes the mean of a linear loss function over the forecast period:

$$\text{MAE} = \frac{1}{N}\sum_{i=1}^{N}\left|\vartheta_i - \widehat{\vartheta}_{j,i}\right| = \frac{1}{N}\sum_{i=1}^{N}|e_{j,i}| = \overline{|e_j|} \tag{7.36}$$

where $\overline{|e_j|}$ represents the MAE of forecast model j. The MAE provides an indication of the overall extent of the deviation of the forecast variance from the realized variance. A larger value for the MAE indicates that the forecast does not perform as well as a forecast with a smaller MAE.

Mean Square Error

The mean square error (MSE) criterion is the most common loss-function based criterion for forecast evaluation. The criterion computes the mean of a quadratic loss function over the forecast period:

$$\text{MSE} = \frac{1}{N}\sum_{i=1}^{N}\left(\vartheta_i - \widehat{\vartheta}_{j,i}\right)^2 = \frac{1}{N}\sum_{i=1}^{N} e_{j,i}^2 = \overline{e_j^2} \tag{7.37}$$

where $\overline{e_j^2}$ represents the MSE of forecast model j. The MSE criterion penalizes large forecast errors proportionately more than the MAE criterion; because the forecast error is squared, large forecast errors are associated with a greater proportionate loss than is the case with absolute forecast errors. In addition, the MSE can be decomposed as:

$$\begin{aligned}
\overline{e_j^2} &= \frac{1}{N}\sum_{i=1}^{N} e_{j,i}^2 \\
&= \frac{1}{N}\sum_{i=1}^{N}\left[(e_{j,i} - \bar{e}_j) + \bar{e}_j\right]^2 \\
&= \frac{1}{N}\sum_{i=1}^{N}(e_{j,i} - \bar{e}_j)^2 + 2\bar{e}_j\frac{1}{N}\sum_{i=1}^{N} e_{j,i} - \bar{e}_j^2 \\
&= \widehat{\text{var}}(e_j) + \bar{e}_j^2
\end{aligned} \tag{7.38}$$

Thus, the MSE represents the sum of the variance of the forecast error and the squared forecast bias.

7.5.2 Regression Tests

In addition to the tests based on loss functions outlined above, there are a number tests which can be conducted in a regression framework to assess directly the relative predictive power of alternative forecast methodologies. Two such regression tests are employed.

Test for Independent Information

The first approach tests whether one model contributes information to the forecast which has predictive power that is independent of the information contributed by a different model. The test comprises a regression of the form:

$$e_{j,t} = \eta_0 + \eta_1 \widehat{\vartheta}_{k,t} + \epsilon_t \tag{7.39}$$

where $e_{j,t}$ represents the forecast error obtained from one of the model under consideration and $\widehat{\vartheta}_{k,t}$ represents the forecast provided by a different model ($j \neq k$). If the forecast from model k contains information that has power in predicting volatility which is independent of that provided by model j, the coefficient η_1 should be significantly different from zero. If, on the other hand, the information contained in forecast k is simply a subset of the information contained within forecast j, η_1 should be zero.

Rather than testing all possible combinations of forecast models, two groups of tests are performed: In the first group, the forecast error from each of the alternative forecasting models is regressed on the forecast from the RD-GARCH model; in the second group, the forecast error from the RD-GARCH model is regressed on the forecasts from the alternative forecasting models. The first group of regressions tests whether the RD-GARCH methodology is sufficient by itself, and the second group tests whether the RD-GARCH methodology provides additional information independent of that provided by the alternative forecasting techniques.

Test for Difference in MSE

A second test which allows for direct comparison of the MSE of alternative forecasts is provided by Ashley et al. (1980). The test defines:

$$\Delta_t = e_{i,t} - e_{j,t} \tag{7.40a}$$
$$\Sigma_t = e_{i,t} + e_{j,t} - \bar{e}_i - \bar{e}_j \tag{7.40b}$$

That is, Δ_t represents the difference between forecast errors for two different forecast models ($i \neq j$) and Σ_t represents the mean-adjusted sum of forecast errors for the same two models. The test comprises a regression of the form:

$$\Delta_t = \eta_0 + \eta_1 \Sigma_t + \epsilon_t \tag{7.41}$$

where, under conventional least-squares assumptions, the intercept can be expressed as:

$$\begin{aligned} \eta_0 &= \bar{\Delta} - \eta_1 \bar{\Sigma} \\ &= \bar{e}_i - \bar{e}_j - \eta_1 \overline{(e_{i,t} + e_{j,t} - (\bar{e}_i + \bar{e}_j))} \\ &= \bar{e}_i - \bar{e}_j - \eta_1 \left(\bar{e}_i + \bar{e}_j - (\bar{e}_i + \bar{e}_j)\right) \\ &= \bar{e}_i - \bar{e}_j \end{aligned} \tag{7.42}$$

which provides a measure of the difference in bias for the two forecast models. If both $\bar{e}_i$ and $\bar{e}_j$ are positive,[5] η_0 significantly positive indicates that forecast i exhibits more bias than forecast j and η_0 significantly negative indicates that forecast j exhibits more bias than forecast i. The slope can be expressed as:

$$\begin{aligned}
\eta_1 &= \frac{\widehat{\text{cov}}\,(\Delta, \Sigma)}{\widehat{\text{var}}\,(\Sigma)} \\
&= \frac{1}{s_\Sigma^2}\overline{\left(\Delta_t - \bar{\Delta}\right)\left(\Sigma_t - \bar{\Sigma}\right)} \\
&= \frac{1}{s_\Sigma^2}\overline{\left(e_{i,t} - e_{j,t} - \bar{e}_i + \bar{e}_j\right)\left(e_{i,t} + e_{j,t} - (\bar{e}_i + \bar{e}_j)\right)} \\
&= \frac{1}{s_\Sigma^2}\overline{\left(e_{i,t}^2 - e_{j,t}^2 - 2\bar{e}_i e_{i,t} + 2\bar{e}_j e_{j,t} + \bar{e}_i\bar{e}_i - \bar{e}_j\bar{e}_j\right)} \\
&= \frac{1}{s_\Sigma^2}\left(\overline{e_i^2} - \overline{e_j^2} - \bar{e}_i\bar{e}_i + \bar{e}_j\bar{e}_j\right) \\
&= \frac{1}{s_\Sigma^2}\left(\widehat{\text{var}}(e_i) - \widehat{\text{var}}(e_j)\right) \qquad (7.43)
\end{aligned}$$

which provides a measure of the difference in forecast error variance for the two models, scaled by $1/s_\Sigma^2$. η_1 significantly positive indicates that the forecast error variance for model i is greater than the forecast error variance for model j (which is to say that model j is preferred to model i). Conversely, η_1 significantly negative indicates that model i is preferred to model j on the basis of forecast error variance.

The direct comparison of MSE between the two forecast models is afforded by the decomposition shown in Equation 7.38. As the MSE can be expressed as the sum of the forecast error variance and the (squared) forecast bias, if one model exhibits both significantly less forecast bias and significantly lower forecast error variance, that model must provide a significantly lower MSE. In terms of the coefficients of Equation 7.41, this means that if both coefficients are of the same sign and both significantly different from zero, the following inferences may be made:[6]

$$\begin{array}{lclcl}
\eta_0, \eta_1 < 0 & \rightarrow & \text{MSE}_i < \text{MSE}_j & \rightarrow & i \succ j \\
\eta_0, \eta_1 > 0 & \rightarrow & \text{MSE}_i > \text{MSE}_j & \rightarrow & i \prec j
\end{array}$$

If one coefficient is significantly different from zero and the other is not, inference regarding MSE is also possible. If one forecast model exhibits significantly less forecast bias or significantly less forecast error variance, and there is no significant difference in the other criterion, the model with the lower bias or error variance must also exhibit a smaller MSE. Thus, the inferences noted above for the situation where both coefficients are significantly different from zero and have the same sign can be extended:

---

[5] If either $\bar{e}_i$ or $\bar{e}_j$ is negative, the forecast error series may be transformed as $e_t' = e_t + 2|\bar{e}|$ which will reverse the sign of the bias without changing the magnitude of the forecast errors.

[6] $i \prec j$ indicates that model j is preferred to model i.

$$\begin{array}{rcccc}
\eta_0, \eta_1 < 0, \text{ either significant} & \rightarrow & \text{MSE}_i < \text{MSE}_j & \rightarrow & i \succ j \\
\eta_0, \eta_1 > 0, \text{ either significant} & \rightarrow & \text{MSE}_i > \text{MSE}_j & \rightarrow & i \prec j \\
\eta_0 < 0,\ \eta_1 > 0,\ \eta_1 \text{ significant} & \rightarrow & \text{MSE}_i > \text{MSE}_j & \rightarrow & i \prec j \\
\eta_0 > 0,\ \eta_1 < 0,\ \eta_1 \text{ significant} & \rightarrow & \text{MSE}_i < \text{MSE}_j & \rightarrow & i \succ j
\end{array}$$

Care must be taken when η_0 is significantly different from zero, η_1 is not significantly different from zero, and the two coefficients have different signs. In this case, it is tempting to make the statistical inference that $\text{MSE}_i > \text{MSE}_j$ when in fact the observed value for $\text{MSE}_i < \text{MSE}_j$. It is not appropriate to conclude that $\text{MSE}_i > \text{MSE}_j$ in this situation, although it is possible to interpret the coefficient which is significantly different from zero by itself.

If the coefficients are both significantly different from zero and of different signs, no inferences regarding MSE are possible. Forecast bias and forecast error variance can be interpreted separately, however.

7.6 Results

As outlined in Section 7.2 above, three different approaches to selecting a calibration period are employed, and for each of these approaches forecasts for both the one-period ahead daily variance and one-period ahead monthly variance constructed. The following section discusses a number of issues surrounding the performance of the different calibration periods, forecast models, and distributional assumptions under which the tests of forecast performance outlined in Section 7.5 are made. The effect of calibration period on parameter estimates is considered first, followed by the performance of daily and monthly variance forecasts made with models calibrated under the assumption of conditionally normally distributed errors, and finally the forecast performance of models calibrated under the assumption of errors from the Student-*t* and generalized error distributions is addressed.

7.6.1 Calibration Periods

The parameter values estimated during the updating calibration period beginning January 2, 1952 are presented in Tables 7.23, 7.24, and 7.25 for the conventional GARCH, Extended GARCH, and RD-GARCH models respectively.[7] The estimated parameters for both the GARCH and Extended GARCH models change very little as additional observations are added to the end of the calibration period, indicating that the choice to update the calibration of the models on a quarterly basis rather than every time a new daily forecast was required is reasonable. The parameter estimates for the RD-GARCH model, however, exhibit considerably more activity, with the values for parameter estimates as well as the set of included parameters changing regularly from calibration period to calibration period. There does appear to be a pattern, however, to the inclusion and exclusion of parameter estimates. With a single exception, whenever estimates for λ are present in all three regime types, estimates for ν are present in only base and low volatility regimes. The reverse is also true, with estimates of λ present in only the base and low volatility regimes when estimates of ν are present in all three regime types. In addition, estimates of β for high volatility regimes are present only when there are estimates of ν for all three regime types. Finally, the single case where there are estimates of β and ν only for the base and low volatility regimes is associated with the only case where there is an estimate of α for high volatility regimes. Given the examples from Chapter 6 of dissimilar sets of parameter estimates generating quite similar NICs and HICs, it is reasonable to suggest that the apparent differences between parameter estimates for the RD-GARCH model are representative of models with similar characteristics but somewhat different numerical representations.

---

[7] In order to maintain readability of the text, these tables are presented at the end of the chapter, beginning on page 165.

The parameter values estimated during the updating calibration period beginning January 2, 1974 are presented in Tables 7.26, 7.27, and 7.28 for the conventional GARCH, Extended GARCH, and RD-GARCH models respectively. The observations made above for the longer calibration period are generally true for the shorter updating calibration period, although there does appear to be slightly more activity in the parameter estimates for both the conventional and Extended GARCH models than was evident with the longer calibration period. The RD-GARCH model also shows considerably more activity, with little evidence of any pattern in the set of included parameters. There are also notable overall differences between the parameter estimates for calibration periods starting January 2, 1974 and those obtained for calibration periods starting January 2, 1952. In spite of the increase in the variability of parameter estimates across the calibration period beginning January 2, 1974, there is sufficient stability to suggest that quarterly re-estimation is reasonable for this calibration approach.

The parameter values estimated during the five year rolling calibration period are presented in Tables 7.29, 7.30, and 7.31 for the conventional GARCH, Extended GARCH, and RD-GARCH models respectively. The conventional GARCH model shows the least variability among parameter estimates, although there are still notable differences from calibration period to calibration period, especially for the parameter α. There is considerably more variability for the Extended GARCH model. The regime dependent portion of the RD-GARCH model appears to have little impact for the short calibration period, with few parameters changing during regimes of differing volatility. This is likely related to the difficulties associated with regime detection in short samples outlined in Section 6.2.1.

Table 7.1: Loss function tests (daily, calibration starting 1952)

Forecast Model	ME	ME[1]	MAE	MSE
Unconditional	9.188305×10^{-6} (0.921)	-2.046249×10^{-6} (−0.452)	7.592404×10^{-5}	5.051068×10^{-8}
5 Day moving	7.502435×10^{-7} (0.072)	-8.513105×10^{-6} (−1.339)	8.432156×10^{-5}	5.574903×10^{-8}
20 Day moving	1.153993×10^{-6} (0.114)	-9.670070×10^{-6} (−1.903)	8.082731×10^{-5}	5.207386×10^{-8}
60 Day moving	6.553853×10^{-6} (0.658)	-4.486378×10^{-6} (−0.963)	7.675492×10^{-5}	5.037223×10^{-8}
Random walk	-1.139556×10^{-8} (−0.001)	-2.811938×10^{-6} (−0.402)	9.778887×10^{-5}	7.762809×10^{-8}
GARCH	6.019634×10^{-6} (0.603)	-4.438550×10^{-6} (−0.891)	7.709811×10^{-5}	5.063765×10^{-8}
Extended GARCH	8.180284×10^{-6} (0.830)	-2.405191×10^{-6} (−0.512)	7.527998×10^{-5}	4.926671×10^{-8}
RD-GARCH	1.554848×10^{-5} (1.581)	4.774582×10^{-6} (1.053)	7.134068×10^{-5}	4.926417×10^{-8}

The table presents the results of the forecast tests based on loss function measures, and covers the 507 one-period ahead daily forecasts made with a calibration period beginning January 2, 1952. ME measures bias in the forecast. MAE and MSE measure the magnitude of the deviation of the forecast variance from the *ex post* realized variance. Numbers in parentheses under the ME figures are t-statistics for a test of no bias (ME = 0).

[1] The ME presented in this column excludes the forecast errors corresponding to the crash and subsequent recovery of October 27 and 28, 1997 and comprises a total of 505 daily one-period ahead forecasts.

7.6.2 Daily Results

Loss Function Tests

Table 7.1 presents the results for the loss function based tests of forecast performance for the one-period ahead daily forecasts constructed with the calibration period beginning January 2, 1952. Tables 7.2 and 7.3 present the results for the same tests for the calibration period beginning January 2, 1974 and the five year rolling calibration period respectively.

Two figures for ME are presented in Table 7.1; the first column of ME figures (the second column of the table) present the mean forecast errors for the complete sample of 507 daily one-period ahead forecasts of return variance, while those in the second column of ME figures (third column of the table) present the mean forecast errors for the 505 daily one-period ahead forecasts which exclude the dates October 27 and 28, 1997. The two sets of figures for ME are presented because the returns associated with the crash and subsequent recovery of October 27 and 28 are uncharacteristically large in magnitude,[8] and it is illustrative to consider the forecast bias produced by models both in the presence and absence of these outliers. The figures presented for the MAE and MSE criteria include these two outliers as the different treatment of large forecast errors is one of the primary motivations for considering both measures.

The *t*-test for forecast bias conducted on the ME criterion fails to reject the null hypothesis of no bias in all cases for both the full 507 forecast sample and the reduced 505 forecast sample. In the full sample, the bias, while not significantly different from zero, has a positive sign in seven of the eight cases[9] indicating that in general the forecast models tend to underestimate the true variance. When the two outliers are removed from the sample, the reverse situation obtains. The bias has a negative sign in seven of the eight cases indicating that the forecast models tend to overestimate the true variance. The inability of the forecast models to predict the high volatility associated with the crash and recovery of October 1997 is associated with a large positive forecast error (ie. the forecast model has dramatically underestimated the true variance). These two large errors have the effect of changing the apparent nature of the forecast bias for six of the eight models. As it is likely not reasonable to expect the forecast models to predict the uncharacteristic events of October 27 and 28, it is more appropriate to consider the ME which excludes these two events to be the appropriate measure of the forecast bias resulting from the various models.

Two of the models do not exhibit this change in sign of the ME criterion. The random walk model uses the variance in the current period as a forecast of the variance for the next period. As a result, the model fails to predict the volatility associated with the crash of October 27 for the same reasons as the other models. The volatility associated with the recovery of October 28, however, will be predicted by the volatility observed during the crash on the previous day. For this reason, the change in the apparent nature of the bias for the random walk model is less notable than for the other models; rather than two unpredictable outliers, there is only one. For the RD-GARCH model, the lack of a change in sign coincidental with the exclusion of the two outlier forecasts is an indication that the model does, in fact, systematically underestimate rather than overestimate the true variance.

The RD-GARCH model provides the forecast with the lowest value for both the MAE and MSE criteria. In both cases, the Extended GARCH model achieves the second lowest value.

The loss function test results for the calibration period beginning January 2, 1974 are presented in Table 7.2, and are similar to those presented in Table 7.1 for the longer updating calibration period. Because the moving average and random walk models make forecasts based upon between one and 60 prior observed returns, the results for the shorter calibration period are identical to the results presented in Table 7.1, and are not repeated in Table 7.2.

The *t*-test for forecast bias conducted on the ME criterion again fails to reject the null hypothesis of no bias for the full 507 forecast sample in all four of the

---

[8] The return on October 27, 1997 was −6.6%, and the return on October 28, the following day, was 3.8%. The next largest return in the forecast period was −2.9% on March 8, 1996.

[9] While none of the individual measures of bias are significantly different from zero, the probability of observing seven of the eight measures with the same sign purely by chance is very small.

Table 7.2: Loss function tests (daily, calibration starting 1974)

Forecast Model	ME	ME[1]	MAE	MSE
Unconditional	-5.845066×10^{-6} (-0.586)	-1.708110×10^{-5} (-3.768)	8.371020×10^{-5}	5.048112×10^{-8}
GARCH	5.367081×10^{-6} (0.540)	-5.297358×10^{-6} (-1.105)	7.683720×10^{-5}	5.017519×10^{-8}
Extended GARCH	6.052341×10^{-6} (0.612)	-4.649668×10^{-6} (-0.994)	7.602677×10^{-5}	4.954565×10^{-8}
RD-GARCH	1.356328×10^{-5} (1.378)	2.680930×10^{-6} (0.597)	7.213379×10^{-5}	4.929272×10^{-8}

The table presents the results of the forecast tests based on loss function measures, and covers the 507 one-period ahead daily forecasts made with a calibration period beginning January 2, 1974. ME measures bias in the forecast. MAE and MSE measure the magnitude of the deviation of the forecast variance from the *ex post* realized variance. Numbers in parentheses under the ME figures are t-statistics for a test of no bias (ME = 0).

[1] The ME presented in this column excludes the forecast errors corresponding to the crash and subsequent recovery of October 27 and 28, 1997 and comprises a total of 505 daily one-period ahead forecasts.

forecast models which have results that are different from those presented in Table 7.1, and fails to reject the null hypothesis of no bias for three of the four models for the reduced 505 forecast sample. The interpretation of these results is the same as the interpretation of the results from the longer updating calibration period: the unconditional variance, GARCH, and Extended GARCH models tend to overestimate the true variance overall, but the large forecast error associated with the crash and subsequent recovery of October 1997 causes an apparent reversal in the nature of the forecast bias for the GARCH and Extended GARCH models. The ME for the unconditional variance model does not change sign, but does become significantly different from zero. Finally, the ME for the RD-GARCH model is again positive but insignificant in both the full and reduced forecast samples, again suggesting that the model systematically underestimates the true variance.

Once again the RD-GARCH model provides the lowest value for both the MAE and MSE criteria, and once again the Extended GARCH model achieves the second lowest value. It is interesting to note that, while the difference in MSE between the Extended GARCH and RD-GARCH models appears first in the fifth significant digit for the longer calibration period, it appears in the third digit for the shorter updating calibration period; the difference in performance (as measured by the MSE) between the two models is greater for the shorter calibration period.

The loss function test results for the five year rolling calibration period are presented in Table 7.3, and are again similar to those presented in Tables 7.1 and 7.2 for the updating calibration periods. As was the case with the shorter updating calibration period, the results for the moving average and random walk models are identical to the results presented in Table 7.1, and are not repeated in Table 7.3.

The t-test for forecast bias conducted on the ME rejects the null hypothesis of

Table 7.3: Loss function tests (daily, five year rolling calibration)

Forecast Model	ME	ME[1]	MAE	MSE
Unconditional	3.768813×10^{-5} (3.784)	2.647136×10^{-5} (5.857)	6.618777×10^{-5}	5.171390×10^{-8}
GARCH	1.945364×10^{-5} (1.962)	8.719668×10^{-6} (1.854)	7.104559×10^{-5}	5.020664×10^{-8}
Extended GARCH	2.423740×10^{-5} (2.463)	1.355610×10^{-5} (2.918)	6.920676×10^{-5}	4.968404×10^{-8}
RD-GARCH	2.439903×10^{-5} (2.481)	1.374340×10^{-5} (2.957)	6.919895×10^{-5}	4.961940×10^{-8}

The table presents the results of the forecast tests based on loss function measures, and covers the 507 one period ahead daily forecasts made with a five year rolling calibration period. ME measures bias in the forecast. MAE and MSE measure the magnitude of the deviation of the forecast variance from the *ex post* realized variance. Numbers in parentheses under the ME figures are t-statistics for a test of no bias (ME = 0).

[1] The ME presented in this column excludes the forecast errors corresponding to the crash and subsequent recovery of October 27 and 28, 1997 and comprises a total of 505 daily one-period ahead forecasts.

no bias in three of the four forecast models which are different due to the different calibration period. This result is the same for both the full 507 forecast sample and the reduced 505 forecast sample. In all cases, the sign of the ME criterion is positive, suggesting that when the five year rolling calibration period is used, all of the GARCH models as well as the unconditional variance model systematically underestimate the true variance. In all cases, exclusion of the observations associated with the crash and subsequent recovery of October 1997 reduces the ME, but in no case is the sign or significance of the measure changed.

The RD-GARCH model again produces the minimum MSE and the Extended GARCH model produces the second lowest value for this measure. The unconditional variance model produces the minimum MAE for this calibration period, however, although the RD-GARCH does produce the second lowest value. With the much shorter calibration period, the unconditional variance model performs slightly better than the RD-GARCH model if the cost associated with forecast errors is linearly related to the magnitude of the errors, although the RD-GARCH model is preferred if the cost associated with large forecast errors is disproportionately larger than the cost associated with small forecast errors.

Tests for Sufficiency of RD-GARCH

Table 7.4 presents results for regression-based tests of the sufficiency of the RD-GARCH methodology for the one-period ahead daily forecasts constructed with the calibration period beginning January 2, 1952. Table 7.5 presents the results for the same tests for the calibration period beginning January 2, 1974 and for the five year rolling calibration period. The tests take the form:

Table 7.4: Tests for sufficiency of RD-GARCH (daily, calibration starting 1952)

Forecast Model	η_0	η_1	R^2
Unconditional	3.561×10^{-3} (2.814×10^{-3})	-57.15 (45.35)	0.003
5 Day moving	9.296×10^{-6} (1.136×10^{-5})	8.872×10^{-2} (8.066×10^{-2})	0.002
20 Day moving	2.488×10^{-5} (1.393×10^{-5})	-1.331×10^{-1} (1.407×10^{-1})	0.002
60 Day moving	2.292×10^{-5} (1.870×10^{-5})	-1.140×10^{-1} (2.458×10^{-1})	0.000
Random walk	4.112×10^{-6} (1.020×10^{-5})	1.605×10^{-1} $(4.329 \times 10^{-2})^*$	0.027
GARCH	2.016×10^{-5} (1.394×10^{-5})	-7.073×10^{-2} (1.513×10^{-1})	0.000
Extended GARCH	1.427×10^{-5} (1.562×10^{-5})	2.031×10^{-2} (1.924×10^{-1})	0.000

The table presents results for regression-based tests of the form: $e_{D,t} = \eta_0 + \eta_1 \widehat{\vartheta}_{j,t} + \epsilon_t$ which test the sufficiency of the RD-GARCH methodology. η_1 significantly different from zero indicates that the forecast from the alternative model ("model j") contains information that has power in predicting volatility which is independent of that provided by the RD-GARCH model. Numbers in parentheses under the parameter estimates represent standard errors. Estimates which are significantly different from zero at the 5% level are indicated with an asterisk. The data set comprises a total of 507 daily one-period ahead forecasts.

$$e_{D,t} = \eta_0 + \eta_1 \widehat{\vartheta}_{j,t} + \epsilon_t \tag{7.44}$$

which test the sufficiency of the RD-GARCH methodology in the sense that η_1 significantly different from zero indicates that the forecast from the alternative model ("model j") contains information that has power in predicting volatility which is independent of that provided by the RD-GARCH model. Evidence that an alternative forecast model contains such information would suggest that the RD-GARCH forecast could be improved upon by a forecast which is a combination of the two methodologies.

While numerically different, the results for all three calibration periods are qualitatively identical. In all cases, η_1 is significantly different from zero for the random walk model, and is not significantly different from zero for any of the remaining models. This indicates that, regardless of the calibration period used, the random walk model is the only model of those tested which provides information that has predictive power for return variance which is independent of that provided by the RD-GARCH model. This suggests that the forecasts provided by the RD-GARCH model could be improved upon by a model which is a combination of the RD-GARCH and random walk models.

Table 7.5: Tests for sufficiency of RD-GARCH (daily, calibration starting 1974 and five year rolling calibration)

Forecast Model	η_0	η_1	R^2
Unconditional	2.376×10^{-3} (1.343×10^{-3})	-30.66 (17.43)	0.006
5 Day moving	3.660×10^{-6} (1.135×10^{-5})	1.405×10^{-1} (8.059×10^{-2})	0.006
20 Day moving	1.548×10^{-5} (1.396×10^{-5})	-2.730×10^{-2} (1.410×10^{-1})	0.000
60 Day moving	1.165×10^{-5} (1.872×10^{-5})	2.959×10^{-2} (2.461×10^{-1})	0.000
Random walk	9.985×10^{-7} (1.018×10^{-5})	1.764×10^{-1} $(4.321 \times 10^{-2})^*$	0.032
GARCH	1.132×10^{-5} (1.555×10^{-5})	3.412×10^{-2} (1.825×10^{-1})	0.000
Extended GARCH	6.307×10^{-6} (1.631×10^{-5})	1.113×10^{-1} (1.995×10^{-1})	0.001

Calibration starting 1974

Forecast Model	η_0	η_1	R^2
Unconditional	7.737×10^{-5} (1.077×10^{-4})	-1.579 (3.199)	0.000
5 Day moving	2.113×10^{-5} (1.137×10^{-5})	4.633×10^{-2} (8.073×10^{-2})	0.001
20 Day moving	3.918×10^{-5} $(1.392 \times 10^{-5})^*$	-2.110×10^{-1} (1.406×10^{-1})	0.004
60 Day moving	4.038×10^{-5} $(1.869 \times 10^{-5})^*$	-2.472×10^{-1} (2.457×10^{-1})	0.002
Random walk	1.375×10^{-5} (1.022×10^{-5})	1.495×10^{-1} $(4.337 \times 10^{-2})^*$	0.023
GARCH	3.861×10^{-5} $(1.466 \times 10^{-5})^*$	-2.745×10^{-1} (2.100×10^{-1})	0.003
Extended GARCH	3.232×10^{-5} $(1.419 \times 10^{-5})^*$	-1.686×10^{-1} (2.173×10^{-1})	0.001

Five year rolling calibration

The table presents results for regression-based tests of the form: $e_{D,t} = \eta_0 + \eta_1 \widehat{\vartheta}_{j,t} + \epsilon_t$ which test the sufficiency of the RD-GARCH methodology. η_1 significantly different from zero indicates that the forecast from the alternative model ("model j") contains information that has power in predicting volatility which is independent of that provided by the RD-GARCH model. Numbers in parentheses under the parameter estimates represent standard errors. Estimates which are significantly different from zero at the 5% level are indicated with an asterisk. The data set comprises a total of 507 daily one-period ahead forecasts.

Table 7.6: Tests for forecast improvement (daily, calibration starting 1952)

Forecast Model	η_0	η_1	R^2
Unconditional	-4.701×10^{-5} $(1.779 \times 10^{-5})^*$	1.009 $(0.266)^*$	0.028
5 Day moving	8.479×10^{-5} $(1.843 \times 10^{-5})^*$	-1.509 $(0.276)^*$	0.056
20 Day moving	3.708×10^{-5} $(1.823 \times 10^{-5})^*$	-6.452×10^{-1} $(2.726 \times 10^{-1})^*$	0.011
60 Day moving	-7.807×10^{-6} (1.800×10^{-5})	2.579×10^{-1} (2.692×10^{-1})	0.002
Random walk	8.633×10^{-5} $(2.190 \times 10^{-5})^*$	-1.551 $(0.327)^*$	0.043
GARCH	4.426×10^{-5} $(1.795 \times 10^{-5})^*$	-6.868×10^{-1} $(2.684 \times 10^{-1})^*$	0.013
Extended GARCH	2.778×10^{-5} (1.779×10^{-5})	-3.521×10^{-1} (2.659×10^{-1})	0.003

The table presents results for regression-based tests of the form: $e_{j,t} = \eta_0 + \eta_1 \widehat{\vartheta}_{D,t} + \epsilon_t$ which test for forecast improvement by the RD-GARCH methodology. η_1 significantly different from zero indicates that the forecast from the RD-GARCH model contains information that has power in predicting volatility which is independent of that provided by the alternative model ("model j"). Numbers in parentheses under the parameter estimates represent standard errors. Estimates which are significantly different from zero at the 5% level are indicated with an asterisk. The data set comprises a total of 507 daily one-period ahead forecasts.

Tests for Forecast Improvement

Table 7.6 presents results for regression-based tests for forecast improvement by the RD-GARCH methodology for the one-period ahead daily forecasts constructed with the calibration period beginning January 2, 1952. Table 7.7 presents the results for the same tests for the calibration period beginning January 2, 1974 and the five year rolling calibration period. The tests take the form:

$$e_{j,t} = \eta_0 + \eta_1 \widehat{\vartheta}_{D,t} + \epsilon_t \tag{7.45}$$

which test for forecast improvement by the RD-GARCH methodology in the sense that η_1 significantly different from zero indicates that the forecast from the RD-GARCH model contains information that has power in predicting volatility which is independent of that provided by the alternative model ("model j"). Evidence that the RD-GARCH model contains such information would suggest that the forecast from the alternative model could be improved upon by a forecast which is a combination of the two methodologies.

For the updating calibration period beginning January 2, 1952, η_1 is significantly different from zero for the unconditional variance model, the five and 20 day moving average models, the random walk model and the conventional GARCH model.

9

Table 7.7: Tests for forecast improvement (daily, calibration starting 1974 and five year rolling calibration)

Forecast Model	η_0	η_1	R^2
Unconditional	-8.004×10^{-5} $(2.197 \times 10^{-5})^*$	1.287 $(0.341)^*$	0.027
5 Day moving	1.117×10^{-4} $(2.276 \times 10^{-5})^*$	-1.924 $(0.353)^*$	0.056
20 Day moving	4.787×10^{-5} $(2.251 \times 10^{-5})^*$	-8.101×10^{-1} $(3.489 \times 10^{-1})^*$	0.011
60 Day moving	-1.230×10^{-5} (2.223×10^{-5})	3.270×10^{-1} (3.446×10^{-1})	0.002
Random walk	1.129×10^{-4} $(2.706 \times 10^{-5})^*$	-1.958 $(0.419)^*$	0.041
GARCH	3.441×10^{-5} (2.216×10^{-5})	-5.037×10^{-1} (3.435×10^{-1})	0.004
Extended GARCH	2.879×10^{-5} (2.204×10^{-5})	-3.943×10^{-1} (3.416×10^{-1})	0.003

Calibration starting 1974

Forecast Model	η_0	η_1	R^2
Unconditional	8.339×10^{-8} (1.421×10^{-5})	8.030×10^{-1} $(2.186 \times 10^{-1})^*$	0.026
5 Day moving	5.557×10^{-5} $(1.477 \times 10^{-5})^*$	-1.171 $(0.227)^*$	0.050
20 Day moving	2.283×10^{-5} (1.459×10^{-5})	-4.629×10^{-1} $(2.245 \times 10^{-1})^*$	0.008
60 Day moving	-4.239×10^{-6} (1.438×10^{-5})	2.305×10^{-1} (2.214×10^{-1})	0.002
Random walk	5.445×10^{-5} $(1.756 \times 10^{-5})^*$	-1.163 $(0.270)^*$	0.035
GARCH	2.596×10^{-5} (1.432×10^{-5})	-1.390×10^{-1} (2.204×10^{-1})	0.001
Extended GARCH	3.167×10^{-5} $(1.421 \times 10^{-5})^*$	-1.587×10^{-1} (2.188×10^{-1})	0.001

Five year rolling calibration

The table presents results for regression-based tests of the form: $e_{j,t} = \eta_0 + \eta_1 \widehat{\vartheta}_{D,t} + \epsilon_t$ which test for forecast improvement by the RD-GARCH methodology. η_1 significantly different from zero indicates that the forecast from the RD-GARCH model contains information that has power in predicting volatility which is independent of that provided by the alternative model ("model j"). Numbers in parentheses under the parameter estimates represent standard errors. Estimates which are significantly different from zero at the 5% level are indicated with an asterisk. The data set comprises a total of 507 daily one-period ahead forecasts.

This indicates that, with the exception of the 60 day moving average and Extended GARCH models, the RD-GARCH model provides information with predictive power which is independent of the information provided by the alternative models. For the unconditional variance model, the five and 20 day moving average models, and the conventional GARCH model, the fact that the RD-GARCH model provides additional independent information with predictive power, coupled with the observation made above that the alternative models do not provide such independent information to the RD-GARCH forecasts, and the fact that the RD-GARCH model outperformed all four of the alternative models based upon the MAE and MSE criteria, suggests that the RD-GARCH is, by itself, superior to these alternative models. As the random walk model and the RD-GARCH models both provide information with predictive power which is independent of the information provided by the other model, the observation made earlier regarding the probable improvement in forecasts which would result from a model combining the random walk and RD-GARCH methodologies is reinforced.

The results for the updating calibration period beginning January 2, 1974 and the five year rolling calibration period are qualitatively identical to each other and similar to those from the longer updating calibration period. For the shorter calibration periods, η_1 is significantly different from zero for the unconditional variance model, the five and 20 day moving average models, and the random walk model. Again, with the exception of the 60 day moving average and two GARCH based models, the RD-GARCH model provides information with predictive power which is independent of the information provided by the alternative models. Once again the results of the tests for forecast improvement, coupled with the results of the tests for sufficiency of the RD-GARCH methodology and the MAE and MSE criteria, suggest that the RD-GARCH is, by itself, superior to the unconditional variance and five and 20 day moving average models for the shorter updating calibration period. For the five year rolling calibration period, the MAE criterion for the unconditional variance model is slightly lower than the value from the RD-GARCH model, so the RD-GARCH model can only be judged superior to the two shorter moving average models in this case. The results from both of the two shorter calibration periods continue to reinforce the observation made earlier that an improvement in forecasts would result from a model combining the random walk and RD-GARCH methodologies.

Tests for Difference in MSE

Table 7.8 presents results for regression-based tests for difference in MSE between the RD-GARCH methodology and the various alternative models for the one-period ahead daily forecasts constructed with the calibration period beginning January 2, 1952. Table 7.9 presents the results for the same tests for the calibration period beginning January 2, 1974 and the five year rolling calibration period. The tests take the form:

$$\Delta_t = \eta_0 + \eta_1 \Sigma_t + \epsilon_t \tag{7.46}$$

Table 7.8: Tests for difference in MSE (daily, calibration starting 1952)

Forecast Model	η_0	η_1	R^2
Unconditional	-6.360×10^{-6} $(1.638 \times 10^{-6})^*$	7.107×10^{-3} (3.685×10^{-3})	0.007
5 Day moving	-1.480×10^{-5} $(4.247 \times 10^{-6})^*$	3.360×10^{-2} $(9.491 \times 10^{-3})^*$	0.024
20 Day moving	-1.439×10^{-5} $(1.817 \times 10^{-6})^*$	1.521×10^{-2} $(4.059 \times 10^{-3})^*$	0.027
60 Day moving	-8.995×10^{-6} $(1.334 \times 10^{-6})^*$	6.607×10^{-3} $(2.999 \times 10^{-3})^*$	0.010
Random walk	-1.554×10^{-5} (8.993×10^{-6})	1.372×10^{-1} $(1.969 \times 10^{-2})^*$	0.088
GARCH	-9.529×10^{-6} $(1.359 \times 10^{-6})^*$	7.963×10^{-3} $(3.053 \times 10^{-3})^*$	0.013
Extended GARCH	-7.368×10^{-6} $(7.038 \times 10^{-7})^*$	9.041×10^{-4} (1.589×10^{-3})	0.001

The table presents results for regression-based tests of the form: $\Delta_t = \eta_0 + \eta_1 \Sigma_t + \epsilon_t$ where $\Delta_t = e_{j,t} - e_{D,t}$ and $\Sigma_t = e_{j,t} + e_{D,t} - \bar{e}_j - \bar{e}_D$ which test for difference in MSE. η_0 and η_1 of the same sign and at least one significantly different from zero indicates that the MSE from one model is significantly better than the MSE from the other model. $\eta_0, \eta_1 > 0$ indicates that $\text{MSE}_D < \text{MSE}_j$; that is the forecast from the RD-GARCH model provides a significantly lower MSE than that provided by the alternative model ("model j"). Numbers in parentheses under the parameter estimates represent standard errors. Estimates which are significantly different from zero at the 5% level are indicated with an asterisk. The data set comprises a total of 507 daily one-period ahead forecasts.

where $\Delta_t = e_{j,t} - e_{D,t}$, $\Sigma_t = e_{j,t} + e_{D,t} - \bar{e}_j - \bar{e}_D$, and the interpretation of the results is as outlined in Section 7.5.2.

In five of the seven tests based on the calibration period beginning in 1952 η_0 and η_1 are both significantly different from zero and are of different signs. This indicates that the two models cannot be ranked based on difference in MSE, although the bias and forecast variance can be interpreted separately. For all five of these models, η_0 is negative and η_1 is positive, indicating that each of these models provides less bias but greater forecast error variance than the RD-GARCH model. As low forecast error variance is likely preferable to low bias, the results indicate that with the longer updating calibration period the RD-GARCH model is preferable to all three of the moving average models, the random walk model, and the conventional GARCH model. The results for the unconditional variance model and the Extended GARCH model should also be interpreted separately since η_0 is significantly different from zero and negative while η_1 is not significantly different from zero but positive. This does indicate that (statistically) the alternative model exhibits significantly less bias than the RD-GARCH model and that there is no difference in forecast error variance between the two models. The observed value of the MSE criterion for both of the alternative models is greater than the observed value for the RD-GARCH model, however. It this situation it appropriate to note that the alternative model

Table 7.9: Tests for difference in MSE (daily, calibration starting 1974 and five year rolling calibration)

Forecast Model	η_0	η_1	R^2
Unconditional	-7.718×10^{-6} $(1.280 \times 10^{-6})^*$	6.749×10^{-3} $(2.875 \times 10^{-3})^*$	0.011
5 Day moving	-1.281×10^{-5} $(4.476 \times 10^{-6})^*$	3.330×10^{-2} $(1.003 \times 10^{-2})^*$	0.021
20 Day moving	-1.241×10^{-5} $(2.073 \times 10^{-6})^*$	1.481×10^{-2} $(4.634 \times 10^{-3})^*$	0.020
60 Day moving	-7.009×10^{-6} $(1.279 \times 10^{-6})^*$	6.163×10^{-3} $(2.873 \times 10^{-3})^*$	0.009
Random walk	-1.355×10^{-5} (9.110×10^{-6})	1.373×10^{-1} $(1.999 \times 10^{-2})^*$	0.085
GARCH	-8.196×10^{-6} $(1.218 \times 10^{-6})^*$	5.247×10^{-3} (2.739×10^{-3})	0.007
Extended GARCH	-7.511×10^{-6} $(9.556 \times 10^{-7})^*$	2.034×10^{-3} (2.154×10^{-3})	0.002

Calibration starting 1974

Forecast Model	η_0	η_1	R^2
Unconditional	1.329×10^{-5} $(1.917 \times 10^{-6})^*$	6.451×10^{-3} (4.322×10^{-3})	0.004
5 Day moving	-2.365×10^{-5} $(4.103 \times 10^{-6})^*$	3.349×10^{-2} $(9.157 \times 10^{-3})^*$	0.026
20 Day moving	-2.325×10^{-5} $(1.759 \times 10^{-6})^*$	1.520×10^{-2} $(3.927 \times 10^{-3})^*$	0.029
60 Day moving	-1.785×10^{-5} $(1.494 \times 10^{-6})^*$	6.606×10^{-3} (3.362×10^{-3})	0.008
Random walk	-2.439×10^{-5} $(8.943 \times 10^{-6})^*$	1.369×10^{-1} $(1.956 \times 10^{-2})^*$	0.088
GARCH	$-4.945 \times 10^{-}$[illegible] $(6.618 \times 10^{-7})^*$	4.072×10^{-3} $(1.489 \times 10^{-3})^*$	0.015
Extended GARCH	-1.616×10^{-7} (9.367×10^{-8})	3.694×10^{-4} (2.115×10^{-4})	0.006

Five year rolling calibration

The table presents results for regression-based tests of the form: $\Delta_t = \eta_0 + \eta_1 \Sigma_t + \epsilon_t$ where $\Delta_t = e_{j,t} - e_{D,t}$ and $\Sigma_t = e_{j,t} + e_{D,t} - \bar{e}_j - \bar{e}_D$ which test for difference in MSE. η_0 and η_1 of the same sign and at least one significantly different from zero indicates that the MSE from one model is significantly better than the MSE from the other model. $\eta_0, \eta_1 > 0$ indicates that $\text{MSE}_D < \text{MSE}_j$; that is the forecast from the RD-GARCH model provides a significantly lower MSE than that provided by the alternative model ("model j"). Numbers in parentheses under the parameter estimates represent standard errors. Estimates which are significantly different from zero at the 5% level are indicated with an asterisk. The data set comprises a total of 507 daily one-period ahead forecasts.

exhibits less bias, but it is not appropriate to conclude that the alternative model provides a lower MSE. Indeed, it is interesting to note that in all cases and for all calibration periods, η_1 is positive, even in the six cases where it is not significantly different from zero. The forecast error variance provided by the RD-GARCH model is less than the forecast error variance provided by the alternative models.

The results for both the shorter updating calibration period and the five year rolling calibration period are qualitatively identical to the results for the longer updating calibration period for all three of the moving average models and the random walk model. In addition, the results for the conventional GARCH model are the same for the longer updating calibration period and the five year rolling calibration period. In all nine cases, η_0 is significantly different from zero and negative and η_1 is significantly different from zero and positive. These results again suggest that, with the shorter calibration periods, the RD-GARCH model is preferable to these alternative forecasting models. The results for the Extended GARCH model is also the same for all three calibration periods with η_0 significantly different from zero and negative and η_1 not significantly different from zero but positive. Once again, it is appropriate to conclude that the Extended GARCH model provides less bias than the RD-GARCH model. It is, however, again not appropriate to conclude that the Extended GARCH model provides a lower MSE than the RD-GARCH model.

The final two models for the shorter updating calibration period are the unconditional volatility model and the conventional GARCH model. For the unconditional volatility model, η_0 is significantly different from zero and negative and η_1 is significantly different from zero and positive, suggesting the same interpretation that RD-GARCH model is preferable. For the conventional GARCH model, η_0 is significantly different from zero and negative and η_1 is not significantly different from zero but positive. Again the observed value of the MSE criterion for the conventional GARCH model is greater than the observed value for the RD-GARCH model.

Finally, the unconditional variance model for the five year rolling calibration period has η_0 and η_1 both positive, with η_0 significantly different from zero. The interpretation in this case is that $\text{MSE}_D < \text{MSE}_u$, which implies that the RD-GARCH model is preferred to the unconditional variance model.

Summary

The results for the daily forecasts based on the three different calibration periods are all similar. In all cases, the RD-GARCH model provides the best fit under the MSE criterion, and in two of the three cases it also provides the best fit under the MAE. In the third case, the RD-GARCH model is second best under the MAE criterion. The MSE for the RD-GARCH model is similar for both updating calibration periods and increases somewhat for the much shorter 5 year rolling calibration period. This is likely because the detection of regimes is problematic in the shorter period, and the resulting regime dependent model is almost identical to the conventional version. For the purposes of forecasting daily variances, the RD-GARCH methodology appears to be quite insensitive to the length of the calibration period, provided the period is long enough that regime detection is effective.

The tests of sufficiency of RD-GARCH and tests for model improvement by RD-GARCH suggest that forecasts of daily variances could be improved by the inclusion of a random walk component in the forecast model. The random walk model provides infinite memory to the forecast in the sense that using yesterday's variance without adjustment as a forecast of today's variance implies that a shock which is impounded into today's variance will pass unadjusted into all future forecasts. Improvement of forecasts due to the inclusion of a random walk process indicates that the RD-GARCH model by itself does not make optimal use of information from the distant past.

The tests for difference in MSE suggest that, with the exception of the unconditional variance model for the rolling calibration period where $\text{MSE}_D < \text{MSE}_u$, there is not a statistically significant difference in MSE among the forecast models. In general, however, the RD-GARCH model provides lower forecast error variance than the alternative models at the expense of greater forecast bias.

Under the presumption that low forecast error variance is preferable to low forecast bias, and as the longer updating calibration period performs slightly better than the shorter updating period with very little extra cost, the RD-GARCH forecast model calibrated with the updating calibration period beginning January 2, 1952 is the preferred model for forecasting daily return variance.

7.6.3 Monthly Results

Loss Function Tests

Table 7.10 presents the results for the loss function based tests of forecast performance for the one-period ahead monthly forecasts constructed with the calibration period beginning January 2, 1952. Tables 7.11 and 7.12 present the results for the same tests for the calibration period beginning January 2, 1974 and the five year rolling calibration period respectively.

As was the case for the daily forecasts, two figures for ME are presented in Table 7.10; the first column of ME figures (the second column of the table) presents the mean forecast errors for the complete sample of 24 monthly one-period ahead forecasts of return variance, while those in the second column of ME figures (third column of the table) present the mean forecast errors for the 23 monthly one-period ahead forecasts which exclude October 1997. Two sets of figures for ME are once again presented because the monthly return variance for October 1997 is uncharacteristically large in magnitude,[10] and it is illustrative to consider the forecast bias produced by models both in the presence and absence of this observation. The figures presented for the MAE and MSE criteria include this forecast as the different treatment of large forecast errors is one of the primary motivations for considering both measures.

The *t*-test for forecast bias conducted on the ME criterion fails to reject the null hypothesis of no bias in seven of the eight cases both for the full 24 forecast sample and for the reduced 23 forecast sample. In both samples, the bias has a

---

[10] The monthly return variance for October 1997 was 6.42×10^{-3}. The next largest return variance in the forecast period was 2.85×10^{-3} for July 1996.

Table 7.10: Loss function tests (monthly, calibration starting 1952)

Forecast Model	ME	ME[1]	MAE	MSE
Unconditional	3.825475×10^{-4} (1.602)	1.819475×10^{-4} (1.283)	7.411279×10^{-4}	1.515423×10^{-6}
5 Day moving	-1.762361×10^{-4} (−0.192)	-4.281250×10^{-4} (−0.464)	1.820172×10^{-3}	2.024397×10^{-5}
20 Day moving	2.646986×10^{-4} (0.837)	5.276963×10^{-5} (0.212)	9.225542×10^{-4}	2.471863×10^{-6}
60 Day moving	3.737920×10^{-4} (1.512)	1.861566×10^{-4} (1.078)	8.163307×10^{-4}	1.606457×10^{-6}
Random walk	6.008489×10^{-4} (1.278)	3.557618×10^{-4} (0.843)	1.480612×10^{-3}	5.668457×10^{-6}
GARCH	2.454211×10^{-4} (0.673)	3.453909×10^{-5} (0.110)	9.524607×10^{-4}	3.254755×10^{-6}
Extended GARCH	4.573449×10^{-4} (1.794)	2.572420×10^{-4} (1.511)	7.535103×10^{-4}	1.768820×10^{-6}
RD-GARCH	5.106392×10^{-4} (2.238)	3.075569×10^{-4} (2.631)	6.626962×10^{-4}	1.510434×10^{-6}

The table presents the results of the forecast tests based on loss function measures, and covers the 24 one-period ahead monthly forecasts made with a calibration period beginning January 2, 1952. ME measures bias in the forecast. MAE and MSE measure the magnitude of the deviation of the forecast variance from the *ex post* realized variance. Numbers in parentheses under the ME figures are t-statistics for a test of no bias (ME $= 0$).

[1] The ME presented in this column excludes the forecast errors corresponding to the crash and subsequent recovery of October 1997 and comprises a total of 23 monthly one-period ahead forecasts.

positive sign in seven of the eight cases[11] indicating that in general the forecast models tend to underestimate the true variance. Unlike the results for the daily forecasts, removing the forecast for October 1997 from the sample does not alter the sign nor the significance of any of the ME measures. The RD-GARCH model is the only model for which the ME of the monthly forecasts of return variance is significantly different from zero. This does not necessarily suggest a weakness in the RD-GARCH model, however, as the RD-GARCH model also provides the forecast with the lowest value for both the MAE and MSE criteria. As was illustrated in the decomposition of the MSE measure presented in Equation 7.38, MSE can be expressed as the sum of the forecast error variance and the squared forecast bias. The statistically significant bias exhibited by the RD-GARCH model, then, indicates that the forecast error variance for the RD-GARCH model is lower than would be suggested by the MSE criterion alone.

The unconditional variance model provides the second lowest value for both the

---

[11] Again, while seven of the eight individual measures of bias are not significantly different from zero, the probability of observing six of the seven measures with the same sign purely by chance is very small.

Table 7.11: Loss function tests (monthly, calibration starting 1974)

Forecast Model	ME	ME[1]	MAE	MSE
Unconditional	6.450326×10^{-5} (0.270)	-1.354036×10^{-4} (−0.947)	7.661879×10^{-4}	1.373679×10^{-6}
GARCH	2.946407×10^{-4} (0.959)	8.737415×10^{-5} (0.363)	8.759080×10^{-4}	2.351933×10^{-6}
Extended GARCH	1.985940×10^{-4} (0.701)	-5.700817×10^{-6} (−0.027)	7.837797×10^{-4}	1.966176×10^{-6}
RD-GARCH	3.987601×10^{-4} (1.642)	1.945290×10^{-4} (1.352)	7.072593×10^{-4}	1.574384×10^{-6}

The table presents the results of the forecast tests based on loss function measures, and covers the 24 one-period ahead monthly forecasts made with a calibration period beginning January 2, 1974. ME measures bias in the forecast. MAE and MSE measure the magnitude of the deviation of the forecast variance from the *ex post* realized variance. Numbers in parentheses under the ME figures are t-statistics for a test of no bias (ME = 0).

[1] The ME presented in this column excludes the forecast errors corresponding to the crash and subsequent recovery of October 1997 and comprises a total of 23 monthly one-period ahead forecasts.

MAE and MSE criteria. The Extended GARCH model, which performed second best for the daily forecasts achieves only the third best value for MAE and the fourth best value for MSE.

The loss function test results for the calibration period beginning January 2, 1974 are presented in Table 7.11. Because the moving average and random walk models make forecasts based upon between one and 60 prior observed returns, the results for the shorter calibration period are identical to the results presented in Table 7.10, and are not repeated in Table 7.11.

The results for the shorter updating calibration period are somewhat different from the results for the longer calibration period. For both the full 24 forecast sample and the reduced 23 forecast sample, the t-test for forecast bias conducted on the ME criterion fails to reject the null hypothesis of no bias for all four of the forecast models which have results that are different due to the shorter calibration period. In two of the cases, the unconditional variance model and the Extended GARCH model, the sign of the ME criterion changes from positive to negative when the forecast associated with October 1997 is excluded, indicating that the effect noted in the daily forecasts is present for the monthly forecasts made with the shorter updating calibration period. The inability of the forecast models to predict the high volatility associated with the crash and recovery of October 1997 is associated with a large positive forecast error (ie. the forecast model has underestimated the true variance). This large error has the effect of changing the apparent nature of the forecast bias for these two models.

The results for the MAE and MSE criteria are also slightly different from the results for the longer updating calibration period. The RD-GARCH model provides the forecast with the lowest value for the MAE criterion and the second lowest value for the MSE criterion. The unconditional variance model provides the second lowest

Table 7.12: Loss function tests (monthly, five year rolling calibration)

Forecast Model	ME	ME[1]	MAE	MSE
Unconditional	9.901528×10^{-4} (4.186)	7.899956×10^{-4} (5.712)	1.033551×10^{-3}	2.323402×10^{-6}
GARCH	6.997201×10^{-4} (2.538)	4.979290×10^{-4} (2.481)	1.013956×10^{-3}	2.313662×10^{-6}
Extended GARCH	4.198662×10^{-4} (1.108)	2.376482×10^{-4} (0.681)	1.183579×10^{-3}	3.621308×10^{-6}
RD-GARCH	7.229273×10^{-4} (3.196)	5.560568×10^{-4} (3.407)	9.389737×10^{-4}	1.750350×10^{-6}

The table presents the results of the forecast tests based on loss function measures, and covers the 24 one-period ahead monthly forecasts made with a five year rolling calibration period. ME measures bias in the forecast. MAE and MSE measure the magnitude of the deviation of the forecast variance from the *ex post* realized variance. Numbers in parentheses under the ME figures are t-statistics for a test of no bias (ME = 0).

[1] The ME presented in this column excludes the forecast errors corresponding to the crash and subsequent recovery of October 1997 and comprises a total of 23 monthly one-period ahead forecasts.

value for the MAE criterion and the lowest value for the MSE criterion. The Extended GARCH model again achieves only the third lowest value for MAE and the fourth lowest value for MSE.

The loss function test results for the five year rolling calibration period are presented in Table 7.12. As was the case with the shorter updating calibration period, the results for the moving average and random walk models are identical to the results presented in Table 7.10, and are not repeated in Table 7.12.

The results for the rolling calibration period are somewhat different from the results for both of the updating calibration periods. For both the full 24 forecast sample and the reduced 23 forecast sample, the t-test for forecast bias conducted on the ME criterion rejects the null hypothesis of no bias for three of the four forecast models which have results that are different due to the shorter calibration period. In all cases, the sign of the ME criterion indicates that the models underestimate the true return variance, regardless of whether the forecast for October 1997 is included or not.

The results for the MAE and MSE criteria are also again somewhat different from the results for the updating calibration periods. The RD-GARCH model no longer provides the forecast with the lowest value for either the MAE or the MSE criteria. The 60 day moving average model provides the lowest value for both the MAE and the MSE, while the 20 day moving average performs second best under the MAE criterion and the RD-GARCH performs second best under the MSE criterion and third best under the MAE criterion.

Table 7.13: Tests for sufficiency of RD-GARCH (monthly, calibration starting 1952)

Forecast Model	η_0	η_1	R^2
Unconditional	-6.521×10^{-3} (4.360×10^{-3})	5.366 (3.322)	0.106
5 Day moving	5.661×10^{-4} $(2.556 \times 10^{-4})^*$	-2.968×10^{-2} (5.174×10^{-2})	0.015
20 Day moving	6.661×10^{-4} (3.505×10^{-4})	-1.089×10^{-1} (1.812×10^{-1})	0.016
60 Day moving	4.660×10^{-4} (4.525×10^{-4})	3.381×10^{-2} (2.916×10^{-1})	0.001
Random walk	5.610×10^{-4} $(2.662 \times 10^{-4})^*$	-4.613×10^{-2} (1.104×10^{-1})	0.008
GARCH	6.104×10^{-4} (3.147×10^{-4})	-6.893×10^{-2} (1.430×10^{-1})	0.010
Extended GARCH	6.091×10^{-4} (4.025×10^{-4})	-7.965×10^{-2} (2.628×10^{-1})	0.004

The table presents results for regression-based tests of the form: $e_{D,t} = \eta_0 + \eta_1 \widehat{\vartheta}_{j,t} + \epsilon_t$ which test the sufficiency of the RD-GARCH methodology. η_1 significantly different from zero indicates that the forecast from the alternative model ("model j") contains information that has power in predicting volatility which is independent of that provided by the RD-GARCH model. Numbers in parentheses under the parameter estimates represent standard errors. Estimates which are significantly different from zero at the 5% level are indicated with an asterisk. The data set comprises a total of 24 monthly one-period ahead forecasts.

Tests for Sufficiency of RD-GARCH

Table 7.13 presents results for regression-based tests of the sufficiency of the RD-GARCH methodology for the one-period ahead monthly forecasts constructed with the calibration period beginning January 2, 1952. Table 7.14 presents the results for the same tests for the calibration period beginning January 2, 1974 and for the five year rolling calibration period.

The results for the two updating calibration periods are qualitatively identical. In both cases, η_1 is not significantly different from zero for any of the alternative forecast models. This indicates that, when calibrated with the updating calibration periods, the RD-GARCH model is sufficient by itself to generate forecasts of return variance; none of the alternative models tested provides information which has predictive power independent of that provided by the RD-GARCH model.

The results are considerably different for the five year updating calibration period. In this case, η_1 is significantly different from zero for the five and 20 day moving average forecast models and for the conventional GARCH model. This suggests that for the much shorter calibration period, the forecasts provided by the RD-GARCH model could be improved upon by a model which is a combination of the RD-GARCH model and at least one of the alternative models which has η_1 significantly different from zero. One interpretation for this observation is that the ability of the RD-GARCH

Table 7.14: Tests for sufficiency of RD-GARCH (monthly, calibration starting 1974 and five year rolling calibration)

Forecast Model	η_0	η_1	R^2	Forecast Model	η_0	η_1	R^2
Unconditional	-5.818×10^{-3} (4.695×10^{-3})	3.817 (2.879)	0.074	Unconditional	-1.170×10^{-3} (2.292×10^{-3})	2.693 (3.243)	0.030
5 Day moving	3.965×10^{-4} (2.740×10^{-4})	1.204×10^{-3} (5.547×10^{-2})	0.000	5 Day moving	9.478×10^{-4} $(2.215 \times 10^{-4})^*$	-1.203×10^{-1} $(4.485 \times 10^{-2})^*$	0.246
20 Day moving	4.184×10^{-4} (3.760×10^{-4})	-1.375×10^{-2} (1.943×10^{-1})	0.000	20 Day moving	1.234×10^{-3} $(3.177 \times 10^{-4})^*$	-3.575×10^{-1} $(1.642 \times 10^{-1})^*$	0.177
60 Day moving	1.287×10^{-4} (4.769×10^{-4})	2.047×10^{-1} (3.073×10^{-1})	0.020	60 Day moving	1.008×10^{-3} $(4.429 \times 10^{-4})^*$	-2.163×10^{-1} (2.854×10^{-1})	0.025
Random walk	5.293×10^{-4} (2.777×10^{-4})	-1.196×10^{-1} (1.151×10^{-1})	0.047	Random walk	8.013×10^{-4} $(2.624 \times 10^{-4})^*$	-7.173×10^{-2} (1.088×10^{-1})	0.019
GARCH	3.331×10^{-4} (3.752×10^{-4})	4.697×10^{-2} (1.979×10^{-1})	0.003	GARCH	1.163×10^{-3} $(2.993 \times 10^{-4})^*$	-4.432×10^{-1} $(2.093 \times 10^{-1})^*$	0.169
Extended GARCH	3.348×10^{-4} (4.180×10^{-4})	4.277×10^{-2} (2.225×10^{-1})	0.002	Extended GARCH	1.021×10^{-3} $(2.642 \times 10^{-4})^*$	-2.337×10^{-1} (1.177×10^{-1})	0.152
Calibration starting 1974				Five year rolling calibration			

The table presents results for regression-based tests of the form: $e_{D,t} = \eta_0 + \eta_1 \hat{\vartheta}_{j,t} + \epsilon_t$ which test the sufficiency of the RD-GARCH methodology. η_1 significantly different from zero indicates that the forecast from the alternative model ("model j") contains information that has power in predicting volatility which is independent of that provided by the RD-GARCH model. Numbers in parentheses under the parameter estimates represent standard errors. Estimates which are significantly different from zero at the 5% level are indicated with an asterisk. The data set comprises a total of 24 monthly one-period ahead forecasts.

Table 7.15: Tests for forecast improvement (monthly, calibration starting 1952)

Forecast Model	η_0	η_1	R^2
Unconditional	-6.216×10^{-4} (6.155×10^{-4})	8.493×10^{-1} (4.817×10^{-1})	0.124
5 Day moving	7.242×10^{-3} $(1.861 \times 10^{-3})^*$	-6.274 $(1.456)^*$	0.458
20 Day moving	1.992×10^{-3} $(7.747 \times 10^{-4})^*$	-1.461 $(0.606)^*$	0.209
60 Day moving	8.023×10^{-4} (6.734×10^{-4})	-3.625×10^{-1} (5.270×10^{-1})	0.021
Random walk	2.514×10^{-3} $(1.217 \times 10^{-3})^*$	-1.618 (0.953)	0.116
GARCH	2.629×10^{-3} $(8.409 \times 10^{-4})^*$	-2.016 $(0.658)^*$	0.299
Extended GARCH	1.476×10^{-3} $(6.614 \times 10^{-4})^*$	-8.614×10^{-1} (5.176×10^{-1})	0.112

The table presents results for regression-based tests of the form: $e_{j,t} = \eta_0 + \eta_1 \widehat{\vartheta}_{D,t} + \epsilon_t$ which test for forecast improvement by the RD-GARCH methodology. η_1 significantly different from zero indicates that the forecast from the RD-GARCH model contains information that has power in predicting volatility which is independent of that provided by the alternative model ("model j"). Numbers in parentheses under the parameter estimates represent standard errors. Estimates which are significantly different from zero at the 5% level are indicated with an asterisk. The data set comprises a total of 24 monthly one-period ahead forecasts.

model to forecast n-period ahead volatility (where $n > 1$) is reduced when the model is calibrated with the short five year rolling calibration period. Alternatively, it is possible that the ability of the conventional GARCH model to forecast n-period ahead volatility improves with shorter calibration periods.

It would make no sense to attempt to combine the RD-GARCH model with a conventional GARCH model, but the two moving average models which have η_1 significantly different from zero indicate that a model which combines the RD-GARCH methodology with an intermediate length moving average component would provide an improved forecast over that provided by the RD-GARCH model by itself for this short calibration period.

Tests for Forecast Improvement

Table 7.15 presents results for regression-based tests for forecast improvement by the RD-GARCH methodology for the one-period ahead monthly forecasts constructed with the calibration period beginning January 2, 1952. Table 7.16 presents the results for the same tests for the calibration period beginning January 2, 1974 and the five year rolling calibration period

For the updating calibration period beginning January 2, 1952, η_1 is significantly different from zero for the five and 20 day moving average models and for the con-

Table 7.16: Tests for forecast improvement (monthly, calibration starting 1974 and five year rolling calibration)

Forecast Model	η_0	η_1	R^2
Unconditional	-5.863×10^{-4} (6.027×10^{-4})	5.028×10^{-1} (4.265×10^{-1})	0.059
5 Day moving	3.315×10^{-3} (2.245×10^{-3})	-2.697 (1.589)	0.116
20 Day moving	1.152×10^{-3} (7.966×10^{-4})	-6.857×10^{-1} (5.637×10^{-1})	0.063
60 Day moving	4.160×10^{-4} (6.431×10^{-4})	-3.261×10^{-2} (4.551×10^{-1})	0.000
Random walk	3.655×10^{-3} $(9.956 \times 10^{-4})^*$	-2.360 $(0.705)^*$	0.338
GARCH	9.230×10^{-4} (7.857×10^{-4})	-4.855×10^{-1} (5.560×10^{-1})	0.033
Extended GARCH	9.091×10^{-4} (7.183×10^{-4})	-5.490×10^{-1} (5.083×10^{-1})	0.050

Calibration starting 1974

Forecast Model	η_0	η_1	R^2
Unconditional	4.140×10^{-4} (3.353×10^{-4})	5.939×10^{-1} $(2.591 \times 10^{-1})^*$	0.193
5 Day moving	3.932×10^{-3} $(8.549 \times 10^{-4})^*$	-4.235 $(0.661)^*$	0.651
20 Day moving	9.337×10^{-4} $(4.614 \times 10^{-4})^*$	-6.896×10^{-1} (3.565×10^{-1})	0.145
60 Day moving	3.628×10^{-4} (3.900×10^{-4})	1.131×10^{-2} (3.014×10^{-1})	0.000
Random walk	9.159×10^{-4} (7.366×10^{-4})	-3.248×10^{-1} (5.691×10^{-1})	0.015
GARCH	1.140×10^{-3} $(4.166 \times 10^{-4})^*$	-4.536×10^{-1} (3.219×10^{-1})	0.083
Extended GARCH	1.780×10^{-3} $(4.556 \times 10^{-4})^*$	-1.402 $(0.352)^*$	0.419

Five year rolling calibration

The table presents results for regression-based tests of the form: $e_{j,t} = \eta_0 + \eta_1 \widehat{\vartheta}_{D,t} + \epsilon_t$ which test for forecast improvement by the RD-GARCH methodology. η_1 significantly different from zero indicates that the forecast from the RD-GARCH model contains information that has power in predicting volatility which is independent of that provided by the alternative model ("model j"). Numbers in parentheses under the parameter estimates represent standard errors. Estimates which are significantly different from zero at the 5% level are indicated with an asterisk. The data set comprises a total of 24 monthly one-period ahead forecasts.

ventional GARCH model. This indicates that, for the longer updating calibration period, the RD-GARCH model provides information with predictive power which is independent of the information provided by these three alternative models. For each of the alternative models, the fact that the RD-GARCH model provides additional independent information with predictive power, coupled with the observation made above that the alternative models do not provide such independent information to the RD-GARCH forecasts, and the fact that the RD-GARCH model outperformed all four of the alternative models based upon the MAE and MSE criteria, suggests that the RD-GARCH is, by itself, superior to the three alternative models.

Unlike the daily results, the results for the updating calibration period beginning January 2, 1974 and for the five year rolling calibration period are different from each other, and are also different from the results for the longer updating calibration period. For the shorter updating calibration period, η_1 is significantly different from zero for only the random walk model, indicating that the RD-GARCH model provides information with predictive power independent of that provided by the random walk model. As the random walk model did not provide information with predictive power independent of that provided by the RD-GARCH model and as the MAE and MSE criteria are both lower for the RD-GARCH model, it is not likely that a model which combines the random walk and RD-GARCH methodologies would provide improved forecasts over those provided by the RD-GARCH model alone.

For the five year rolling calibration period, η_1 is significantly different from zero for the unconditional variance, five day moving average, and Extended GARCH models. As the five day moving average and RD-GARCH models both provide information with predictive power which is independent of the information provided by the other model, it is likely that the forecast made with the RD-GARCH model could be improved upon by a model which combined a relatively short moving average with the RD-GARCH methodology. This reinforces the observation made in the previous section that an increase in forecast performance would probably result from the inclusion of a moving average component with the RD-GARCH model.

Tests for Difference in MSE

Table 7.17 presents results for regression-based tests for difference in MSE between the RD-GARCH methodology and the various alternative models for the one-period ahead monthly forecasts constructed with the calibration period beginning January 2, 1952. Table 7.18 presents the results for the same tests for the calibration period beginning January 2, 1974 and for the five year rolling calibration period.

In four of the seven tests based on the calibration period beginning January 2, 1952, the null hypothesis of no difference in MSE can be rejected in favour of the alternative that $\text{MSE}_D < \text{MSE}_j$. For the five and 20 day moving average model, the random walk model, and the conventional GARCH model, η_1 is significantly different from zero and positive and η_0 is not significantly different from zero. The forecast error variance is greater for the alternative model than for the RD-GARCH model, and there is no statistically significant difference in bias.

The shorter updating calibration period beginning January 2, 1974 provides two

Table 7.17: Tests for difference in MSE (monthly, calibration starting 1952)

Forecast Model	η_0	η_1	R^2
Unconditional	-1.281×10^{-4} (1.034×10^{-4})	2.388×10^{-2} (4.625×10^{-2})	0.012
5 Day moving	-3.344×10^{-4} (3.979×10^{-4})	7.545×10^{-1} $(7.938 \times 10^{-2})^*$	0.804
20 Day moving	-2.459×10^{-4} (1.712×10^{-4})	1.786×10^{-1} $(6.741 \times 10^{-2})^*$	0.242
60 Day moving	-1.368×10^{-4} (1.210×10^{-4})	4.255×10^{-2} (5.359×10^{-2})	0.028
Random walk	9.021×10^{-5} (2.978×10^{-4})	4.310×10^{-1} $(9.706 \times 10^{-2})^*$	0.473
GARCH	-2.652×10^{-4} (2.250×10^{-4})	2.681×10^{-1} $(8.354 \times 10^{-2})^*$	0.319
Extended GARCH	-5.329×10^{-5} (1.024×10^{-4})	5.772×10^{-2} (4.421×10^{-2})	0.072

The table presents results for regression-based tests of the form: $\Delta_t = \eta_0 + \eta_1 \Sigma_t + \epsilon_t$ where $\Delta_t = e_{j,t} - e_{D,t}$ and $\Sigma_t = e_{j,t} + e_{D,t} - \bar{e}_j - \bar{e}_D$ which test for difference in MSE. η_0 and η_1 of the same sign and at least one significantly different from zero indicates that the MSE from one model is significantly better than the MSE from the other model. $\eta_0, \eta_1 > 0$ indicates that $\text{MSE}_D < \text{MSE}_j$; that is the forecast from the RD-GARCH model provides a significantly lower MSE than that provided by the alternative model ("model j"). Numbers in parentheses under the parameter estimates represent standard errors. Estimates which are significantly different from zero at the 5% level are indicated with an asterisk. The data set comprises a total of 24 monthly one-period ahead forecasts.

of the seven alternative models with η_1 significantly different from zero and positive and η_0 not significantly different from zero. Again, for the five day moving average and random walk models the null hypothesis of no difference in MSE can be rejected in favour of the alternative that $\text{MSE}_D < \text{MSE}_j$. The unconditional variance model has η_0 significantly different from zero and negative, while and η_1 is not significantly different from zero but also negative, indicating that in this case the null hypothesis of no difference in MSE can be rejected in favour of the alternative that $\text{MSE}_u < \text{MSE}_D$; this indicates that the difference in MSE noted in the loss function based tests presented in Table 7.14 is statistically significant.

The five year rolling calibration period provides four of the seven alternative models with η_1 significantly different from zero and positive and η_0 not significantly different from zero. For the five day moving average, random walk, and both GARCH based models the null hypothesis of no difference in MSE can be rejected in favour of the alternative that $\text{MSE}_D < \text{MSE}_j$. For the 60 day moving average variance model, η_0 is significantly different from zero and negative while η_1 is not significantly different from zero but positive. This indicates that (statistically) the alternative model exhibits significantly less bias than the RD-GARCH model and that there is no significant difference in forecast error variance between the two models. For the 20 day moving average variance model, η_0 is significantly different from zero and negative

Table 7.18: Tests for difference in MSE (monthly, calibration starting 1974 and five year rolling calibration)

Forecast Model	η_0	η_1	R^2
Unconditional	-3.343×10^{-4} $(1.226 \times 10^{-4})^*$	-8.753×10^{-3} (5.355×10^{-2})	0.001
5 Day moving	-2.225×10^{-4} (4.523×10^{-4})	7.804×10^{-1} $(9.215 \times 10^{-2})^*$	0.765
20 Day moving	-1.341×10^{-4} (2.245×10^{-4})	1.548×10^{-1} (8.893×10^{-2})	0.121
60 Day moving	-2.497×10^{-5} (1.714×10^{-4})	1.004×10^{-2} (7.575×10^{-2})	0.001
Random walk	2.021×10^{-4} (2.732×10^{-4})	3.765×10^{-1} $(8.498 \times 10^{-2})^*$	0.472
GARCH	-1.041×10^{-4} (2.341×10^{-4})	1.408×10^{-1} (9.530×10^{-2})	0.090
Extended GARCH	-2.002×10^{-4} (2.021×10^{-4})	8.909×10^{-2} (8.436×10^{-2})	0.048

Calibration starting 1974

Forecast Model	η_0	η_1	R^2
Unconditional	2.672×10^{-4} (1.766×10^{-4})	2.589×10^{-2} (8.371×10^{-2})	0.004
5 Day moving	-5.467×10^{-4} (2.861×10^{-4})	6.690×10^{-1} $(5.370 \times 10^{-2})^*$	0.876
20 Day moving	-4.582×10^{-4} $(1.165 \times 10^{-4})^*$	1.738×10^{-1} $(4.484 \times 10^{-2})^*$	0.406
60 Day moving	-3.491×10^{-4} $(1.485 \times 10^{-4})^*$	4.885×10^{-2} (6.712×10^{-2})	0.024
Random walk	-1.221×10^{-4} (3.115×10^{-4})	4.479×10^{-1} $(1.032 \times 10^{-1})^*$	0.461
GARCH	-2.321×10^{-5} (8.100×10^{-5})	1.011×10^{-1} $(3.335 \times 10^{-2})^*$	0.295
Extended GARCH	-3.031×10^{-4} (1.605×10^{-4})	2.711×10^{-1} $(5.611 \times 10^{-2})^*$	0.515

Five year rolling calibration

The table presents results for regression-based tests of the form: $\Delta_t = \eta_0 + \eta_1 \Sigma_t + \epsilon_t$ where $\Delta_t = e_{j,t} - e_{D,t}$ and $\Sigma_t = e_{j,t} + e_{D,t} - \bar{e}_j - \bar{e}_D$ which test for difference in MSE. η_0 and η_1 of the same sign and at least one significantly different from zero indicates that the MSE from one model is significantly better than the MSE from the other model. $\eta_0, \eta_1 > 0$ indicates that $\mathrm{MSE}_D < \mathrm{MSE}_j$; that is the forecast from the RD-GARCH model provides a significantly lower MSE than that provided by the alternative model ("model j"). Numbers in parentheses under the parameter estimates represent standard errors. Estimates which are significantly different from zero at the 5% level are indicated with an asterisk. The data set comprises a total of 24 monthly one-period ahead forecasts.

while η_1 is significantly different from zero and positive; the moving average model provides less bias but greater forecast error variance than the RD-GARCH model. Under the presumption noted earlier that forecasts with low forecast error variance are preferable to forecasts with low bias, the RD-GARCH model is preferred over the 20 day moving average model.

Finally, it is interesting to note that in all cases and for all calibration periods, with the single exception of the unconditional variance model for the shorter updating calibration period (which was the preferred model under the MSE criterion for this calibration period), η_1 is positive—even in nine of the ten cases where it is not significantly different from zero. It is reasonable to conclude from this that the forecast error variance provided by the RD-GARCH model is in general less than the forecast error variance provided by the alternative models.

Summary

The results for monthly forecasts exhibit considerably less consistency across the calibration periods than the results for the daily forecasts. The minimum MSE and MAE are again provided by the RD-GARCH methodology for the longer updating calibration period, with different models preferred for both shorter periods. For the models constructed with the shorter updating calibration period, the unconditional variance model is preferred by the MSE criterion, although the RD-GARCH model does provide the minimum MAE and the second best MSE. For the five year rolling calibration period, the 60 day moving average variance model is preferred by both the MAE and the MSE criteria. The RD-GARCH model provides the second best MSE, but only the third best MAE, with the 20 day moving average model achieving the second lowest value. Overall, the RD-GARCH model provides the lowest MAE for two of the three calibration periods, and the lowest MSE for the longer updating calibration period. The preferred model in the remaining cases is different for each calibration period. The RD-GARCH model calibrated with the longer updating calibration period also provides the lowest overall MAE and the second lowest MSE with only the unconditional variance model calibrated with the shorter updating period performing better.

The tests of sufficiency of RD-GARCH and tests for model improvement by RD-GARCH suggest that for both updating calibration periods, the forecasts of monthly variance provided by the RD-GARCH model would not be improved by a model combining the RD-GARCH methodology with any of the other models. For the five year rolling average calibration period, the tests suggest that the forecasts from the RD-GARCH model could be improved by the inclusion a component based on a moving average of past variances. As the longer updating calibration periods perform better overall than the short rolling calibration period, however, it is more appropriate to use the longer calibration period rather than attempting to construct such a composite model.

The tests for difference in MSE indicate that in ten of the 21 cases RD-GARCH provides a MSE which is significantly less than that provided by the alternative model. There is one case where the reverse it true: for the shorter updating calibration period

$\text{MSE}_u < \text{MSE}_D$. With the exception of the unconditional variance model for the shorter updating calibration period, the RD-GARCH model provides lower forecast error variance than the alternative models at the expense of greater forecast bias.

The RD-GARCH model provides the greatest consistency in the forecasting of monthly return variance. It is the preferred model for the longer updating calibration period, and is the second best model for both of the shorter calibration periods. The preferred model for the shorter updating calibration period is the unconditional variance model, while the preferred model for the five year rolling calibration period is the 60 day moving average model. The consistency provided by the RD-GARCH model is useful because, as outlined in Section 7.2, the selection of the calibration period is ad-hoc in nature. The problem of calibration period selection is particularly evident in forecasting monthly variance as the preferred model is different for each calibration period. In light of this issue, and because of its overall performance, the RD-GARCH forecast model calibrated with the updating calibration period beginning January 2, 1952 is again the preferred model for forecasting monthly return variance.

7.6.4 Alternative Distributional Assumptions

Sections 4.4 and 5.5 developed and tested versions of the RD-GARCH model under two alternative distributions for the disturbances: the generalized error distribution (GED) and the Student-t distribution. While both alternative distributions provided better model fit, as measured by likelihood values, than the version of the model constructed under the assumption of normally distributed disturbances, the results of the sign bias tests suggested that the assumption of normality for disturbances provides greater independence for the standardized residuals than did either of the alternative distributions. This section addresses the lack of a clear preference for one of the three distributional assumptions by assessing the relative performance of the alternatives at out-of-sample forecasting.

Tables 7.19 and 7.20 present the results for the loss function based tests of forecast performance and the regression based tests for difference in MSE for the one-period ahead daily forecasts constructed with the calibration period beginning January 2, 1952.

None of the daily forecast models exhibit significant bias as measured by the ME criterion, although the model constructed under the assumption of normally distributed disturbances exhibits least bias and the model constructed under the assumption of disturbances from the GED produces forecasts with the greatest bias. The model constructed under the assumption of disturbances from the GED provides the preferred forecast under the MAE criterion, and the model constructed under the assumption of Student-t distributed disturbances provides the preferred model under the MSE criterion.

For the GED, Table 7.20 shows η_0 significantly different from zero and positive, and η_1 not significantly different from zero but positive. The model constructed under the assumption of normally distributed disturbances outperforms the model constructed under the assumption of disturbances from the GED based on the MSE criterion. For the Student-t distribution, η_0 is significantly different from zero and

Table 7.19: Loss function tests (daily, alternative distributions)

Error Distribution	ME	ME[1]	MAE	MSE
Normal	1.554848×10^{-5} (1.581)	4.774582×10^{-6} (1.053)	7.134068×10^{-5}	4.926417×10^{-8}
GED	1.710547×10^{-5} (1.739)	6.257003×10^{-6} (1.390)	7.069414×10^{-5}	4.932047×10^{-8}
Student-*t*	1.617637×10^{-5} (1.654)	5.629700×10^{-6} (1.241)	7.096042×10^{-5}	4.874291×10^{-8}

The table presents the results of the forecast tests based on loss function measures, and covers the 507 one-period ahead daily forecasts made with a calibration period beginning January 2, 1952. ME measures bias in the forecast. MAE and MSE measure the magnitude of the deviation of the forecast variance from the *ex post* realized variance. Numbers in parentheses under the ME figures are t-statistics for a test of no bias (ME = 0).

[1] The ME presented in this column excludes the forecast errors corresponding to the crash and subsequent recovery of October 27 and 28, 1997 and comprises a total of 505 daily one-period ahead forecasts.

Table 7.20: Tests for difference in MSE (daily, alternative distributions)

Error Distribution	η_0	η_1	R^2
GED	1.557×10^{-6} $(1.905 \times 10^{-7})^*$	2.784×10^{-5} (4.303×10^{-4})	0.000
Student-*t*	6.279×10^{-7} $(2.513 \times 10^{-7})^*$	-2.776×10^{-3} $(5.692 \times 10^{-4})^*$	0.045

The table presents results for regression-based tests of the form: $\Delta_t = \eta_0 + \eta_1 \Sigma_t + \epsilon_t$ where $\Delta_t = e_{j,t} - e_{D,t}$ and $\Sigma_t = e_{j,t} + e_{D,t} - \bar{e}_j - \bar{e}_D$ which test for difference in MSE. η_0 and η_1 of the same sign and at least one significantly different from zero indicates that the MSE from one model is significantly better than the MSE from the other model. $\eta_0, \eta_1 > 0$ indicates that $\text{MSE}_N < \text{MSE}_j$; that is the forecast made from the model which assumes normally distributed errors provides a significantly lower MSE than that provided by the model which makes an alternative distributional assumption ("model j"). Numbers in parentheses under the parameter estimates represent standard errors. Estimates which are significantly different from zero at the 5% level are indicated with an asterisk. The data set comprises a total of 507 daily one-period ahead forecasts.

positive, and η_1 is significantly different from zero and negative, requiring that the parameters be interpreted separately. $\eta_0 > 0$ indicates that the model constructed under the assumption of normally distributed disturbances exhibits less bias than the model constructed under the assumption of Student-t disturbances, while $\eta_1 < 0$ indicates that the model constructed under the assumption of Student-t disturbances exhibits less forecast error variance than the model constructed under the assumption of normally distributed disturbances.

Under the presumption that low forecast error variance is preferable to low forecast bias, and noting that $\text{MSE}_t < \text{MSE}_N$, this suggests that the assumption of Student-t disturbances provides forecasts which are superior to those provided under the

Table 7.21: Loss function tests (monthly, alternative distributions)

Error Distribution	ME	ME[1]	MAE	MSE
Normal	5.106392×10^{-4} (2.238)	3.075569×10^{-4} (2.631)	6.626962×10^{-4}	1.510434×10^{-6}
GED	5.835010×10^{-4} (2.574)	3.831319×10^{-4} (3.230)	7.109264×10^{-4}	1.573914×10^{-6}
Student-*t*	4.845117×10^{-4} (1.967)	2.805140×10^{-4} (1.864)	7.628362×10^{-4}	1.691348×10^{-6}

The table presents the results of the forecast tests based on loss function measures, and covers the 23 one-period ahead monthly forecasts made with a calibration period beginning January 2, 1952. ME measures bias in the forecast. MAE and MSE measure the magnitude of the deviation of the forecast variance from the *ex post* realized variance. Numbers in parentheses under the ME figures are *t*-statistics for a test of no bias (ME = 0).

[1] The ME presented in this column excludes the forecast errors corresponding to the crash and subsequent recovery of October 1997 and comprises a total of 23 monthly one-period ahead forecasts.

Table 7.22: Tests for difference in MSE (monthly, alternative distributions)

Error Distribution	η_0	η_1	R^2
GED	7.286×10^{-5} $(2.139 \times 10^{-5})^*$	-3.277×10^{-3} (9.608×10^{-3})	0.005
Student-*t*	-2.613×10^{-5} (7.735×10^{-5})	3.924×10^{-2} (3.369×10^{-2})	0.058

The table presents results for regression-based tests of the form: $\Delta_t = \eta_0 + \eta_1 \Sigma_t + \epsilon_t$ where $\Delta_t = e_{j,t} - e_{D,t}$ and $\Sigma_t = e_{j,t} + e_{D,t} - \bar{e}_j - \bar{e}_D$ which test for difference in MSE. η_0 and η_1 of the same sign and at least one significantly different from zero indicates that the MSE from one model is significantly better than the MSE from the other model. $\eta_0, \eta_1 > 0$ indicates that $\text{MSE}_N < \text{MSE}_j$; that is the forecast made from the model which assumes normally distributed errors provides a significantly lower MSE than that provided by the model which makes an alternative distributional assumption ("model *j*"). Numbers in parentheses under the parameter estimates represent standard errors. Estimates which are significantly different from zero at the 5% level are indicated with an asterisk. The data set comprises a total of 24 monthly one-period ahead forecasts.

assumption of normally distributed disturbances, and also provides forecasts which are superior to those provided by the GED.

Tables 7.21 and 7.22 present the results for the loss function based tests of forecast performance and the regression based tests for difference in MSE for the one-period ahead monthly forecasts constructed with the calibration period beginning January 2, 1952.

Both the model constructed under the assumption of normally distributed disturbances and the model constructed under the assumption of disturbances from the GED exhibit significant bias as measured by the ME criterion, while the model constructed under the assumption of Student-*t* disturbances does not. The model constructed under the assumption of normally distributed disturbances provides the preferred forecast under both the MAE and the MSE criteria.

For the GED, Table 7.22 shows η_0 significantly different from zero and positive, and η_1 not significantly different from zero but negative. As the computed values for MSE show $\text{MSE}_N < \text{MSE}_{GED}$, the inference may be made that the model constructed under the assumption of normally distributed disturbances outperforms the model constructed under the assumption of disturbances from the GED based on the MSE criterion. For the Student-t distribution, neither η_0 nor η_1 is significantly different from zero. However, as the model constructed under the assumption of Student-t disturbances exhibits lower bias and higher values for the MSE criterion than does the model constructed under the assumption of normally distrusted disturbances, it is reasonable to conclude that $\text{MSE}_N < \text{MSE}_t$. For forecasts of monthly return variance, the model constructed under the assumption of normally distributed errors provides the superior forecast.

Table 7.23: Calibration parameters (GARCH, starting 1952)

Ending	γ_0	γ_2	γ_3	γ_4	ω	α	β
December 29, 1995	$8.450^{[-4]}$ $(7.988^{[-5]})^*$	0.212 $(0.011)^*$	−0.045 $(0.010)^*$	−0.002 $(1.565^{[-4]})^*$	$8.306^{[-7]}$ $(1.894^{[-7]})^*$	0.093 $(0.016)^*$	0.896 $(0.016)^*$
March 29, 1996	$8.402^{[-4]}$ $(7.926^{[-5]})^*$	0.211 $(0.011)^*$	−0.044 $(0.010)^*$	−0.002 $(1.559^{[-4]})^*$	$8.327^{[-7]}$ $(1.898^{[-7]})^*$	0.093 $(0.016)^*$	0.897 $(0.016)^*$
June 28, 1996	$8.343^{[-4]}$ $(7.836^{[-5]})^*$	0.211 $(0.011)^*$	−0.044 $(0.010)^*$	−0.002 $(1.553^{[-4]})^*$	$8.246^{[-7]}$ $(1.861^{[-7]})^*$	0.092 $(0.016)^*$	0.897 $(0.016)^*$
September 30, 1996	$8.357^{[-4]}$ $(7.820^{[-5]})^*$	0.211 $(0.011)^*$	−0.044 $(0.010)^*$	−0.002 $(1.550^{[-4]})^*$	$8.456^{[-7]}$ $(1.905^{[-7]})^*$	0.093 $(0.016)^*$	0.896 $(0.016)^*$
December 31, 1996	$8.395^{[-4]}$ $(7.827^{[-5]})^*$	0.212 $(0.011)^*$	−0.043 $(0.010)^*$	−0.002 $(1.545^{[-4]})^*$	$8.476^{[-7]}$ $(1.895^{[-7]})^*$	0.093 $(0.016)^*$	0.896 $(0.016)^*$
March 31, 1997	$8.393^{[-4]}$ $(7.816^{[-5]})^*$	0.213 $(0.011)^*$	−0.044 $(0.010)^*$	−0.002 $(1.542^{[-4]})^*$	$8.425^{[-7]}$ $(1.879^{[-7]})^*$	0.093 $(0.016)^*$	0.897 $(0.016)^*$
June 30, 1997	$8.307^{[-4]}$ $(7.026^{[-5]})^*$	0.210 $(0.011)^*$	−0.044 $(0.010)^*$	−0.002 $(1.506^{[-4]})^*$	$7.648^{[-7]}$ $(1.974^{[-7]})^*$	0.090 $(0.016)^*$	0.901 $(0.016)^*$
September 30, 1997	$8.444^{[-4]}$ $(7.785^{[-5]})^*$	0.210 $(0.011)^*$	−0.045 $(0.010)^*$	−0.002 $(1.538^{[-4]})^*$	$8.308^{[-7]}$ $(1.851^{[-7]})^*$	0.092 $(0.016)^*$	0.898 $(0.016)^*$

Entries of the form $x^{[-y]}$ represent scientific notation; that is $2.5^{[-2]}$ represents the value 2.5×10^{-2}.

The table presents the parameter estimates which result when the conventional GARCH model is estimated for the calibration period which begins January 2, 1952 and ends as indicated in the first column of the table. The model parameters are then used to forecast the one-period ahead conditional variance for the trading days which begin 1 day following the date indicated in the first column, and end on the date indicated in the first column of the next row of the table. The final date for which a one-period ahead daily forecast is made is December 31, 1997. Forecasts of one-period ahead monthly variances are made from the parameter estimates by summing the individual one-period ahead through N_m-period ahead daily forecasts, where N_m represents the number of trading days in the relevant month. Numbers in parentheses under the parameter estimates represent standard errors. Parameters which are significantly different from zero at the 5% level are indicated with an asterisk.

Table 7.24: Calibration parameters (extended GARCH, starting 1952)

Ending	γ_0	γ_2	γ_3	γ_4	λ	ν	b	ω	α	β
December 29, 1995	$6.227^{[-4]}$ $(6.343^{[-5]})^*$	0.214 (0.012)*	−0.039 (0.013)*	−0.002 $(1.589^{[-4]})^*$	0.247 (0.504)	1.192 (0.166)*	0.512 (0.128)*	0.006 (0.016)	0.019 (0.039)	0.961 (0.038)*
March 29, 1996	$6.200^{[-4]}$ $(6.382^{[-5]})^*$	0.213 (0.012)*	−0.038 (0.016)*	−0.001 $(1.548^{[-4]})^*$	0.228 (0.371)	1.182 (0.138)*	0.509 (0.141)*	0.007 (0.013)	0.018 (0.029)	0.962 (0.029)*
June 28, 1996	$6.161^{[-4]}$ $(6.257^{[-5]})^*$	0.213 (0.012)*	−0.038 (0.011)*	−0.001 $(1.561^{[-4]})^*$	0.235 (0.493)	1.189 (0.162)*	0.513 (0.094)*	0.007 (0.016)	0.018 (0.038)	0.962 (0.037)*
September 30, 1996	$6.140^{[-4]}$ $(6.131^{[-5]})^*$	0.213 (0.011)*	−0.038 (0.010)*	−0.001 $(1.466^{[-4]})^*$	0.230 (0.454)	1.186 (0.154)*	0.517 (0.064)*	0.007 (0.015)	0.018 (0.035)	0.962 (0.035)*
December 31, 1996	$6.164^{[-4]}$ $(6.391^{[-5]})^*$	0.214 (0.011)*	−0.038 (0.011)*	−0.001 $(1.747^{[-4]})^*$	0.225 (0.719)	1.184 (0.214)*	0.519 (0.092)*	0.007 (0.025)	0.018 (0.055)	0.962 (0.052)*
March 31, 1997	$6.158^{[-4]}$ $(6.154^{[-5]})^*$	0.214 (0.011)*	−0.038 (0.010)*	−0.001 $(1.528^{[-4]})^*$	0.212 (0.668)	1.181 (0.201)*	0.520 (0.059)*	0.008 (0.025)	0.017 (0.052)	0.963 (0.049)*
June 30, 1997	$6.206^{[-4]}$ $(6.131^{[-5]})^*$	0.214 (0.011)*	−0.038 (0.010)*	−0.001 $(1.452^{[-4]})^*$	0.022 (0.361)	1.136 (0.126)*	0.519 (0.052)*	0.020 (0.033)	0.002 (0.029)	0.977 (0.029)*
September 30, 1997	$6.214^{[-4]}$ $(6.401^{[-5]})^*$	0.212 (0.011)*	−0.039 (0.014)*	−0.001 $(1.763^{[-4]})^*$	0.194 (0.408)	1.179 (0.145)*	0.517 (0.129)*	0.008 (0.016)	0.015 (0.032)	0.965 (0.031)*

Entries of the form $x^{[-y]}$ represent scientific notation; that is $2.5^{[-2]}$ represents the value 2.5×10^{-2}.

The table presents the parameter estimates which result when the Extended GARCH model is estimated for the calibration period which begins January 2, 1952 and ends as indicated in the first column of the table. The model parameters are then used to forecast the one-period ahead conditional variance for the trading days which begin 1 day following the date indicated in the first column, and end on the date indicated in the first column of the next row of the table. The final date for which a one-period ahead daily forecast is made is December 31, 1997. Forecasts of one-period ahead monthly variances are made from the parameter estimates by summing the individual one-period ahead through N_m-period ahead daily forecasts, where N_m represents the number of trading days in the relevant month. Numbers in parentheses under the parameter estimates represent standard errors. Parameters which are significantly different from zero at the 5% level are indicated with an asterisk.

Table 7.25: Calibration parameters (RD-GARCH, starting 1952)

Ending	γ_0	γ_2	γ_3	γ_4	λ	ν	b	ω	α	β
December 29, 1995	$6.760^{[-4]}$	0.212	−0.036	−0.001	5.212	2.369	0.799	$2.010^{[-13]}$	0.117	0.743
	$(6.387^{[-5]})^*$	(0.011)*	(0.010)*	$(1.474^{[-4]})^*$	(1.716)*	(0.369)*	(0.095)*	$(1.741^{[-12]})$	(0.032)*	(0.067)*
					−3.998	−0.859	0	$1.464^{[-4]}$	0	0
					(1.455)*	(0.336)*		$(4.371^{[-4]})$		
					−2.053	0	0	$3.195^{[-8]}$	0	0
					(1.271)			$(1.603^{[-7]})$		
March 29, 1996	$6.694^{[-4]}$	0.211	−0.035	−0.001	4.682	2.238	0.806	$2.704^{[-12]}$	0.111	0.769
	$(6.418^{[-5]})^*$	(0.011)*	(0.010)*	$(1.577^{[-4]})^*$	(1.515)*	(0.332)*	(0.101)*	$(2.077^{[-11]})$	(0.033)*	(0.061)*
					−3.591	−0.824	0	$2.544^{[-4]}$	0	0
					(1.349)*	(0.324)*		$(6.224^{[-4]})$		
					0	0.863	0	$2.872^{[-11]}$	0	−0.107
						(0.584)		$(2.015^{[-10]})$		(0.104)
June 28, 1996	$6.663^{[-4]}$	0.211	−0.034	−0.001	4.887	2.264	0.811	$9.200^{[-13]}$	0.111	0.768
	$(6.327^{[-5]})^*$	(0.011)*	(0.010)*	$(1.496^{[-4]})^*$	(0.746)*	(0.271)*	(0.092)*	$(3.418^{[-12]})$	(0.036)*	(0.057)*
					−3.799	−0.852	0	$2.618^{[-4]}$	0	0
					(0.754)*	(0.380)*		$(7.553^{[-4]})$		
					0	0.913	0	$1.092^{[-11]}$	0	−0.115
						(0.569)		$(3.840^{[-11]})$		(0.101)
September 30, 1996	$6.686^{[-4]}$	0.211	−0.034	−0.001	5.388	2.367	0.811	$8.500^{[-14]}$	0.122	0.734
	$(6.358^{[-5]})^*$	(0.011)*	(0.010)*	$(1.488^{[-4]})^*$	(2.098)*	(0.389)*	(0.094)*	$(9.002^{[-13]})$	(0.034)*	(0.077)*
					−4.071	−0.806	0	$8.322^{[-5]}$	0	0
					(1.703)*	(0.326)*		$(2.962^{[-4]})$		
					−2.172	0	0	$2.486^{[-8]}$	0	0
					(1.456)			$(1.370^{[-7]})$		
December 31, 1996	$6.684^{[-4]}$	0.212	−0.034	−0.001	5.624	2.429	0.819	$2.518^{[-14]}$	0.120	0.728
	$(6.370^{[-5]})^*$	(0.011)*	(0.010)*	$(1.460^{[-4]})^*$	(1.419)*	(0.308)*	(0.090)*	$(1.767^{[-13]})$	(0.052)*	(0.089)*
					−4.306	−0.843	0	$8.551^{[-5]}$	0	0
					(0.947)*	(0.693)		$(6.243^{[-4]})$		
					−2.270	0	0	$1.296^{[-8]}$	0	0
					(1.184)			$(6.033^{[-8]})$		

Entries of the form $x^{[-y]}$ represent scientific notation; that is $2.5^{[-2]}$ represents the value 2.5×10^{-2}.

Table 7.25: continued.

Ending	γ_0	γ_2	γ_3	γ_4	λ	ν	b	ω	α	β
March 31, 1997	$6.672^{[-4]}$	0.213	−0.035	−0.001	5.638	2.376	0.828	$2.080^{[-14]}$	0.120	0.739
	$(6.285^{[-5]})$*	(0.011)*	(0.010)*	$(1.460^{[-4]})$*	(0.127)*	(0.236)*	(0.082)*	$(1.331^{[-14]})$	(0.022)*	(0.039)*
					−4.323	−0.823	0	$8.016^{[-5]}$	0	0
					(0.122)*	(0.132)*		$(9.310^{[-9]})$*		
					0	0.961	0	$3.191^{[-13]}$	0	−0.128
						$(8.160^{[-4]})$*		$(7.660^{[-18]})$*		$(2.366^{[-6]})$*
June 30, 1997	$6.709^{[-4]}$	0.212	−0.035	−0.001	4.669	2.263	0.814	$2.839^{[-12]}$	0.107	0.773
	$(6.285^{[-5]})$*	(0.011)*	(0.010)*	$(1.478^{[-4]})$*	(1.371)*	(0.322)*	(0.096)*	$(1.971^{[-11]})$	(0.031)*	(0.057)*
					−3.633	−0.868	0	$3.464^{[-4]}$	0	0
					(1.244)*	(0.323)*		$(7.731^{[-4]})$		
					0	0.869	0	$3.036^{[-11]}$	0	−0.110
						(0.581)		$(1.928^{[-10]})$		(0.104)
September 30, 1997	$6.694^{[-4]}$	0.211	−0.036	−0.001	3.892	2.363	0.808	$1.439^{[-10]}$	0.078	0.812
	$(6.279^{[-5]})$*	(0.011)*	(0.010)*	$(1.467^{[-4]})$*	(1.276)*	(0.345)*	(0.087)*	$(9.326^{[-10]})$	(0.028)*	(0.053)*
					−3.114	−0.987	0	0.001	0	0
					(1.154)*	(0.367)*		(0.003)		
					0	0	0	$1.113^{[-9]}$	0.094	−0.158
								$(6.601^{[-9]})$	(0.055)	(0.104)

Entries of the form $x^{[-y]}$ represent scientific notation; that is $2.5^{[-2]}$ represents the value 2.5×10^{-2}.

The table presents the parameter estimates which result when the RD-GARCH model is estimated for the calibration period which begins January 2, 1952 and ends as indicated in the first column of the table. The model parameters are then used to forecast the one-period ahead conditional variance for the trading days which begin 1 day following the date indicated in the first column, and end on the date indicated in the first column of the next row of the table. The final date for which a one-period ahead daily forecast is made is December 31, 1997. Forecasts of one-period ahead monthly variances are made from the parameter estimates by summing the individual one-period ahead through N_m-period ahead daily forecasts, where N_m represents the number of trading days in the relevant month. Numbers in parentheses under the parameter estimates represent standard errors. Parameters which are significantly different from zero at the 5% level are indicated with an asterisk. The first pair of rows in each section presents the parameter estimates and standard errors for the base regime, and the remaining four rows present the parameter estimates and standard errors for regimes of unusually low and high volatility respectively. Parameters which are excluded by the stepwise estimation procedure are represented by a value of zero, with no corresponding standard error presented.

Table 7.26: Calibration parameters (GARCH, starting 1974)

Ending	γ_0	γ_2	γ_3	γ_4	ω	α	β
December 29, 1995	$6.136^{[-4]}$ $(1.148^{[-4]})^*$	0.160 (0.014)*	−0.021 (0.014)	$-5.431^{[-4]}$ $(2.532^{[-4]})^*$	$9.283^{[-7]}$ $(4.683^{[-7]})$	0.066 (0.025)*	0.922 (0.027)*
March 29, 1996	$6.056^{[-4]}$ $(1.146^{[-4]})^*$	0.159 (0.014)*	−0.020 (0.014)	$-5.004^{[-4]}$ $(2.524^{[-4]})$	$9.465^{[-7]}$ $(4.699^{[-7]})^*$	0.066 (0.025)*	0.922 (0.027)*
June 28, 1996	$5.971^{[-4]}$ $(1.128^{[-4]})^*$	0.159 (0.014)*	−0.020 (0.014)	$-4.731^{[-4]}$ $(2.506^{[-4]})$	$9.276^{[-7]}$ $(4.498^{[-7]})^*$	0.066 (0.025)*	0.923 (0.026)*
September 30, 1996	$6.025^{[-4]}$ $(1.120^{[-4]})^*$	0.161 (0.014)*	−0.020 (0.014)	$-4.983^{[-4]}$ $(2.481^{[-4]})^*$	$9.892^{[-7]}$ $(4.805^{[-7]})^*$	0.067 (0.025)*	0.921 (0.027)*
December 31, 1996	$6.130^{[-4]}$ $(1.112^{[-4]})^*$	0.163 (0.014)*	−0.019 (0.014)	$-4.943^{[-4]}$ $(2.453^{[-4]})^*$	$9.841^{[-7]}$ $(4.613^{[-7]})^*$	0.067 (0.025)*	0.920 (0.027)*
March 31, 1997	$6.141^{[-4]}$ $(1.108^{[-4]})^*$	0.164 (0.014)*	−0.020 (0.014)	$-5.035^{[-4]}$ $(2.440^{[-4]})^*$	$9.801^{[-7]}$ $(4.561^{[-7]})^*$	0.067 (0.024)*	0.921 (0.026)*
June 30, 1997	$6.263^{[-4]}$ $(1.104^{[-4]})^*$	0.163 (0.014)*	−0.022 (0.014)	$-5.050^{[-4]}$ $(2.439^{[-4]})^*$	$9.790^{[-7]}$ $(4.557^{[-7]})^*$	0.066 (0.024)*	0.921 (0.026)*
September 30, 1997	$6.309^{[-4]}$ $(1.105^{[-4]})^*$	0.161 (0.014)*	−0.023 (0.014)	$-4.980^{[-4]}$ $(2.433^{[-4]})^*$	$9.653^{[-7]}$ $(4.490^{[-7]})^*$	0.066 (0.024)*	0.922 (0.026)*

Entries of the form $x^{[-y]}$ represent scientific notation; that is $2.5^{[-2]}$ represents the value 2.5×10^{-2}.

The table presents the parameter estimates which result when the conventional GARCH model is estimated for the calibration period which begins January 2, 1974 and ends as indicated in the first column of the table. The model parameters are then used to forecast the one-period ahead conditional variance for the trading days which begin 1 day following the date indicated in the first column, and end on the date indicated in the first column of the next row of the table. The final date for which a one-period ahead daily forecast is made is December 31, 1997. Forecasts of one-period ahead monthly variances are made from the parameter estimates by summing the individual one-period ahead through N_m-period ahead daily forecasts, where N_m represents the number of trading days in the relevant month. Numbers in parentheses under the parameter estimates represent standard errors. Parameters which are significantly different from zero at the 5% level are indicated with an asterisk.

Table 7.27: Calibration parameters (extended GARCH, starting 1974)

Ending	γ_0	γ_2	γ_3	γ_4	λ	ν	b	ω	α	β
December 31, 1995	$4.586^{[-4]}$ $(1.057^{[-4]})^*$	0.161 $(0.014)^*$	−0.012 (0.014)	$-4.686^{[-4]}$ $(2.461^{[-4]})$	2.826 (2.107)	2.027 $(0.735)^*$	0.503 $(0.093)^*$	$1.718^{[-8]}$ $(1.738^{[-7]})$	0.091 (0.052)	0.878 $(0.091)^*$
March 29, 1996	$4.514^{[-4]}$ $(1.053^{[-4]})^*$	0.160 $(0.014)^*$	−0.011 (0.014)	$-4.261^{[-4]}$ $(2.441^{[-4]})$	2.828 (2.258)	1.993 $(0.736)^*$	0.504 $(0.098)^*$	$1.782^{[-8]}$ $(1.931^{[-7]})$	0.095 (0.060)	0.874 $(0.099)^*$
June 28, 1996	$4.451^{[-4]}$ $(1.046^{[-4]})^*$	0.160 $(0.014)^*$	−0.010 (0.014)	$-4.018^{[-4]}$ $(2.427^{[-4]})$	2.883 (1.738)	2.007 $(0.624)^*$	0.502 $(0.096)^*$	$1.317^{[-8]}$ $(1.092^{[-7]})$	0.095 (0.048)	0.874 $(0.082)^*$
September 30, 1996	$4.443^{[-4]}$ $(1.039^{[-4]})^*$	0.162 $(0.014)^*$	−0.010 (0.014)	$-4.173^{[-4]}$ $(2.408^{[-4]})$	2.982 (2.127)	2.035 $(0.719)^*$	0.518 $(0.096)^*$	$8.595^{[-9]}$ $(8.748^{[-8]})$	0.098 (0.054)	0.867 $(0.096)^*$
December 31, 1996	$4.526^{[-4]}$ $(1.033^{[-4]})^*$	0.163 $(0.014)^*$	−0.010 (0.014)	$-4.087^{[-4]}$ $(2.379^{[-4]})$	2.974 (1.681)	2.032 $(0.626)^*$	0.526 $(0.095)^*$	$8.908^{[-9]}$ $(7.138^{[-8]})$	0.098 $(0.045)^*$	0.867 $(0.081)^*$
March 31, 1997	$4.515^{[-4]}$ $(1.028^{[-4]})^*$	0.165 $(0.014)^*$	−0.011 (0.014)	$-4.122^{[-4]}$ $(2.363^{[-4]})$	2.881 (2.269)	2.011 $(0.753)^*$	0.529 $(0.096)^*$	$1.420^{[-8]}$ $(1.551^{[-7]})$	0.096 (0.055)	0.870 $(0.098)^*$
June 30, 1997	$4.666^{[-4]}$ $(1.023^{[-4]})^*$	0.164 $(0.014)^*$	−0.013 (0.014)	$-4.219^{[-4]}$ $(2.365^{[-4]})$	2.790 (2.136)	1.995 $(0.721)^*$	0.523 $(0.093)^*$	$2.198^{[-8]}$ $(2.258^{[-7]})$	0.093 (0.053)	0.874 $(0.092)^*$
September 30, 1997	$4.690^{[-4]}$ $(1.022^{[-4]})^*$	0.162 $(0.014)^*$	−0.014 (0.014)	$-4.112^{[-4]}$ $(2.356^{[-4]})$	2.765 (1.998)	1.996 $(0.697)^*$	0.528 $(0.096)^*$	$2.492^{[-8]}$ $(2.388^{[-7]})$	0.091 (0.049)	0.876 $(0.087)^*$

Entries of the form $x^{[-y]}$ represent scientific notation; that is $2.5^{[-2]}$ represents the value 2.5×10^{-2}.

The table presents the parameter estimates which result when the Extended GARCH model is estimated for the calibration period which begins January 2, 1974 and ends as indicated in the first column of the table. The model parameters are then used to forecast the one-period ahead conditional variance for the trading days which begin 1 day following the date indicated in the first column, and end on the date indicated in the first column of the next row of the table. The final date for which a one-period ahead daily forecast is made is December 31, 1997. Forecasts of one-period ahead monthly variances are made from the parameter estimates by summing the individual one-period ahead through N_m-period ahead daily forecasts, where N_m represents the number of trading days in the relevant month. Numbers in parentheses under the parameter estimates represent standard errors. Parameters which are significantly different from zero at the 5% level are indicated with an asterisk.

Table 7.28: Calibration parameters (RD-GARCH, starting 1974)

Ending	γ_0	γ_2	γ_3	γ_4	λ	ν	b	ω	α	β
December 29, 1995	$4.955^{[-4]}$ $(1.016^{[-4]})^*$	0.161 (0.015)*	−0.016 (0.014)	$-4.505^{[-4]}$ $(2.411^{[-4]})$	3.123 (1.221)*	1.989 (0.365)*	0.641 (0.061)*	$1.130^{[-8]}$ $(7.542^{[-8]})$	0.065 (0.034)	0.863 (0.001)*
					0	0	2.778 (0.001)*	0	0	−0.863 $(1.597^{[-6]})^*$
					0	0	0	0	0.161 $(1.634^{[-5]})^*$	−0.125 $(6.842^{[-7]})^*$
March 29, 1996	$4.529^{[-4]}$ $(1.024^{[-4]})^*$	0.162 (0.014)*	−0.014 (0.014)	$-3.342^{[-4]}$ $(2.397^{[-4]})$	3.158 (0.058)*	2.022 (0.048)*	0.629 (0.050)*	$1.091^{[-8]}$ $(2.912^{[-9]})^*$	0.071 $(3.958^{[-5]})^*$	0.849 $(1.951^{[-5]})^*$
					0	0	2.550 (0.002)*	0	0	−0.849 $(3.381^{[-5]})^*$
					0	0	0	0	0.163 $(1.049^{[-5]})^*$	−0.119 $(3.127^{[-7]})^*$
June 28, 1996	$4.444^{[-4]}$ $(1.008^{[-4]})^*$	0.167 (0.015)*	−0.011 (0.015)	$-2.677^{[-4]}$ $(2.385^{[-4]})$	2.857 (0.256)*	2.272 (0.265)*	0.716 (0.309)*	$2.277^{[-8]}$ $(5.855^{[-8]})$	0.032 (0.012)*	0.917 (0.051)*
					0	0	3.082 (0.515)*	$4.448^{[-8]}$ $(4.024^{[-10]})^*$	0	−0.917 $(1.963^{[-4]})^*$
					0	0	0	$9.553^{[-8]}$ $(1.563^{[-12]})^*$	0.142 $(1.252^{[-6]})^*$	−0.201 $(6.340^{[-7]})^*$
September 30, 1996	$4.820^{[-4]}$ $(1.001^{[-4]})^*$	0.160 (0.014)*	−0.015 (0.014)	$-4.068^{[-4]}$ $(2.365^{[-4]})$	2.946 (0.148)*	2.140 (0.289)*	0.604 (0.112)*	$1.513^{[-8]}$ $(1.430^{[-8]})$	0.042 (0.017)*	0.911 (0.028)*
					0	0	3.119 (0.985)*	$2.352^{[-8]}$ $(2.188^{[-8]})$	0	−0.911 (0.261)*
					0	0	0	$4.148^{[-8]}$ $(1.331^{[-11]})^*$	0.158 $(2.175^{[-6]})^*$	−0.170 $(3.465^{[-7]})^*$
December 31, 1996	$5.211^{[-4]}$ $(9.953^{[-5]})^*$	0.164 (0.014)*	−0.014 (0.014)	$-3.581^{[-4]}$ $(2.344^{[-4]})$	9.555 (0.103)*	3.618 (0.528)*	0.761 (0.147)*	$2.609^{[-22]}$ $(3.805^{[-23]})^*$	0.073 (0.036)*	0.668 (0.083)*
					−0.759 (0.107)*	0	0	0	0	0
					0	0	0	0	0.233 (0.167)	−0.205 (0.153)

Entries of the form $x^{[-y]}$ represent scientific notation; that is $2.5^{[-2]}$ represents the value 2.5×10^{-2}.

Table 7.28: continued.

Ending	γ_0	γ_2	γ_3	γ_4	λ	ν	b	ω	α	β
March 31, 1997	$4.860^{[-4]}$	0.170	−0.013	$-3.298^{[-4]}$	10.919	5.475	1.922	$2.174^{[-25]}$	0.001	0.728
	$(1.003^{[-4]})^*$	(0.014)*	(0.013)	$(2.392^{[-4]})$	(2.918)*	(0.907)*	(0.100)*	$(3.146^{[-24]})$	(0.001)	(0.061)*
					−7.058	−4.186	0	$2.716^{[-10]}$	0.317	−0.728
					(2.915)*	(0.997)*		$(2.418^{[-11]})^*$	(0.167)	(0.183)*
					−9.167	−4.234	−1.331	$1.708^{[-5]}$	0.257	0
					(2.823)*	(0.795)*	(0.019)*	$(1.886^{[-10]})^*$	$(7.451^{[-7]})^*$	
June 30, 1997	$5.125^{[-4]}$	0.164	−0.018	$-3.419^{[-4]}$	9.185	2.554	0.672	$1.182^{[-21]}$	0.170	0.681
	$(9.950^{[-5]})^*$	(0.014)*	(0.014)	$(2.332^{[-4]})$	(0.110)*	(0.482)*	(0.166)*	$(4.913^{[-22]})^*$	(0.020)*	(0.068)*
					−7.088	0	1.744	$2.074^{[-6]}$	−0.158	0
					(0.104)*		(1.517)	$(7.311^{[-10]})^*$	$(6.078^{[-9]})^*$	
					−5.744	0	0	$1.156^{[-8]}$	0	0
					$(4.511^{[-6]})^*$			$(1.864^{[-27]})^*$		
September 30, 1997	$5.208^{[-4]}$	0.164	−0.017	$-3.230^{[-4]}$	3.379	2.054	0.641	$3.593^{[-9]}$	0.072	0.846
	$(1.046^{[-4]})^*$	(0.016)*	(0.013)	$(2.306^{[-4]})$	(1.853)	(0.486)*	(0.069)*	$(3.623^{[-8]})$	(0.051)	$(6.630^{[-5]})^*$
					0	0	2.402	0	0	−0.846
							$(2.469^{[-4]})^*$			$(7.763^{[-6]})^*$
					0	0	0	0	0.161	−0.121
									$(8.691^{[-6]})^*$	$(2.844^{[-7]})^*$

Entries of the form $x^{[-y]}$ represent scientific notation; that is $2.5^{[-2]}$ represents the value 2.5×10^{-2}.

The table presents the parameter estimates which result when the RD-GARCH model is estimated for the calibration period which begins January 2, 1974 and ends as indicated in the first column of the table. The model parameters are then used to forecast the one-period ahead conditional variance for the trading days which begin 1 day following the date indicated in the first column, and end on the date indicated in the first column of the next row of the table. The final date for which a one-period ahead daily forecast is made is December 31, 1997. Forecasts of one-period ahead monthly variances are made from the parameter estimates by summing the individual one-period ahead through N_m-period ahead daily forecasts, where N_m represents the number of trading days in the relevant month. Numbers in parentheses under the parameter estimates represent standard errors. Parameters which are significantly different from zero at the 5% level are indicated with an asterisk. The first pair of rows in each section presents the parameter estimates and standard errors for the base regime, and the remaining four rows present the parameter estimates and standard errors for regimes of unusually low and high volatility respectively. Parameters which are excluded by the stepwise estimation procedure are represented by a value of zero, with no corresponding standard error presented.

Table 7.29: Calibration parameters (GARCH, five year rolling)

Ending	γ_0	γ_2	γ_3	γ_4	ω	α	β
December 29, 1995	$3.591^{[-4]}$ $(1.776^{[-4]})^*$	0.094 $(0.030)^*$	0.015 (0.031)	$7.169^{[-4]}$ $(3.967^{[-4]})$	$1.268^{[-7]}$ $(4.426^{[-8]})^*$	0	0.995 $(1.673^{[-7]})^*$
March 29, 1996	$3.221^{[-4]}$ $(1.817^{[-4]})$	0.105 $(0.030)^*$	−0.012 (0.032)	$8.854^{[-4]}$ $(4.122^{[-4]})^*$	$4.095^{[-7]}$ $(5.092^{[-7]})$	0.021 (0.020)	0.966 $(0.035)^*$
June 29, 1996	$3.160^{[-4]}$ $(1.805^{[-4]})$	0.118 $(0.033)^*$	−0.011 (0.034)	$9.463^{[-4]}$ $(4.108^{[-4]})^*$	$6.257^{[-7]}$ $(1.130^{[-6]})$	0.026 (0.036)	0.953 $(0.070)^*$
September 30, 1996	$3.298^{[-4]}$ $(1.731^{[-4]})$	0.123 $(0.030)^*$	$-7.109^{[-4]}$ (0.031)	$8.675^{[-4]}$ $(3.921^{[-4]})^*$	$6.686^{[-7]}$ $(5.399^{[-7]})$	0.027 (0.020)	0.949 $(0.035)^*$
December 31, 1996	$4.088^{[-4]}$ $(1.678^{[-4]})^*$	0.132 $(0.030)^*$	−0.009 (0.031)	$6.573^{[-4]}$ $(3.810^{[-4]})$	$1.605^{[-6]}$ $(1.083^{[-6]})$	0.054 $(0.022)^*$	0.892 $(0.052)^*$
March 31, 1997	$4.497^{[-4]}$ $(1.682^{[-4]})^*$	0.157 $(0.031)^*$	−0.010 (0.030)	$5.681^{[-4]}$ $(3.857^{[-4]})$	$1.392^{[-6]}$ $(8.053^{[-7]})$	0.058 $(0.019)^*$	0.897 $(0.038)^*$
June 30, 1997	$5.266^{[-4]}$ $(1.697^{[-4]})^*$	0.151 $(0.031)^*$	−0.015 (0.030)	$4.059^{[-4]}$ $(3.870^{[-4]})$	$9.224^{[-7]}$ $(5.325^{[-7]})$	0.057 $(0.018)^*$	0.916 $(0.030)^*$
September 30, 1997	$5.773^{[-4]}$ $(1.754^{[-4]})^*$	0.143 $(0.030)^*$	−0.022 (0.031)	$3.895^{[-4]}$ $(3.906^{[-4]})$	$7.115^{[-7]}$ $(4.371^{[-7]})$	0.054 $(0.018)^*$	0.927 $(0.027)^*$

Entries of the form $x^{[-y]}$ represent scientific notation; that is $2.5^{[-2]}$ represents the value 2.5×10^{-2}.

The table presents the parameter estimates which result when the conventional GARCH model is estimated for a five year rolling calibration period. In each case, the calibration period encompasses the five calendar years period which ends as indicated in the first column of the table. The model parameters are then used to forecast the one-period ahead conditional variance for the trading days which begin 1 day following the date indicated in the first column, and end on the date indicated in the first column of the next row of the table. The final date for which a one-period ahead daily forecast is made is December 31, 1997. Forecasts of one-period ahead monthly variances are made from the parameter estimates by summing the individual one-period ahead through N_m-period ahead daily forecasts, where N_m represents the number of trading days in the relevant month. Numbers in parentheses under the parameter estimates represent standard errors. Parameters which are significantly different from zero at the 5% level are indicated with an asterisk.

Table 7.30: Calibration parameters (extended GARCH, five year rolling)

Ending	γ_0	γ_2	γ_3	γ_4	λ	ν	b	ω	α	β
December 29, 1995	$3.641^{[-4]}$	0.101	0.002	$6.938^{[-4]}$	0.062	4.881	1.455	0.003	$1.122^{[-6]}$	0.995
	$(1.724^{[-4]})^*$	$(0.030)^*$	(0.031)	$(3.966^{[-4]})$	(0.396)	$(1.260)^*$	(2.804)	(0.005)	$(8.313^{[-6]})$	$(0.004)^*$
March 29, 1996	$2.246^{[-4]}$	0.121	0.013	$7.538^{[-4]}$	0.381	1.091	0.798	0.014	0.034	0.865
	$(1.798^{[-4]})$	$(0.036)^*$	(0.034)	$(5.038^{[-4]})$	(8.482)	(1.347)	$(0.103)^*$	(0.604)	(0.718)	(0.748)
June 28, 1996	$1.766^{[-4]}$	0.143	0.023	$9.028^{[-4]}$	$1.255^{[-6]}$	1.011	1.105	0.208	$1.480^{[-7]}$	0.792
	$(1.694^{[-4]})$	$(0.033)^*$	(0.037)	$(4.104^{[-4]})^*$	$(7.704^{[-7]})$	$(0.390)^*$	$(0.095)^*$	$(8.721^{[-7]})^*$	$(1.368^{[-11]})^*$	$(1.134^{[-20]})^*$
September 30, 1996	$1.824^{[-4]}$	0.138	0.028	$8.824^{[-4]}$	2.121	1.569	1.048	$1.284^{[-6]}$	0.094	0.776
	$(2.034^{[-4]})$	$(0.036)^*$	(0.046)	$(4.179^{[-4]})^*$	(1.341)	$(0.609)^*$	(0.541)	$(7.238^{[-6]})$	(0.195)	(0.433)
December 31, 1996	$2.121^{[-4]}$	0.138	0.005	$6.868^{[-4]}$	3.646	2.659	2.208	$5.153^{[-10]}$	0.031	0.572
	$(1.665^{[-4]})$	$(0.030)^*$	(0.031)	$(3.806^{[-4]})$	$(1.638)^*$	$(0.480)^*$	$(0.810)^*$	$(4.445^{[-9]})$	(0.025)	$(0.193)^*$
March 31, 1997	$2.318^{[-4]}$	0.167	0.006	$6.720^{[-4]}$	3.676	2.563	1.705	$4.213^{[-10]}$	0.052	0.609
	$(1.715^{[-4]})$	$(0.031)^*$	(0.032)	$(3.770^{[-4]})$	$(1.610)^*$	$(0.618)^*$	$(0.754)^*$	$(3.513^{[-9]})$	(0.067)	$(0.199)^*$
June 30, 1997	$3.953^{[-4]}$	0.166	−0.010	$3.601^{[-4]}$	$1.379^{[-6]}$	1.474	0.727	0.063	$7.424^{[-8]}$	0.937
	$(1.679^{[-4]})^*$	$(0.031)^*$	(0.031)	$(3.731^{[-4]})$	$(5.593^{[-7]})^*$	$(0.259)^*$	$(0.138)^*$	$(1.979^{[-7]})^*$	$(2.221^{[-13]})^*$	$(5.404^{[-22]})^*$
September 30, 1997	$4.033^{[-4]}$	0.163	−0.006	$4.388^{[-4]}$	6.252	2.724	0.834	$2.738^{[-16]}$	0.158	0.655
	$(1.872^{[-4]})^*$	$(0.035)^*$	(0.034)	$(3.915^{[-4]})$	(3.886)	$(0.666)^*$	$(0.364)^*$	$(5.733^{[-15]})$	(0.156)	$(0.205)^*$

Entries of the form $x^{[-y]}$ represent scientific notation; that is $2.5^{[-2]}$ represents the value 2.5×10^{-2}.

The table presents the parameter estimates which result when the Extended GARCH model is estimated for a five year rolling calibration period. In each case, the calibration period encompasses the five calendar years period which ends as indicated in the first column of the table. The model parameters are then used to forecast the one-period ahead conditional variance for the trading days which begin 1 day following the date indicated in the first column, and end on the date indicated in the first column of the next row of the table. The final date for which a one-period ahead daily forecast is made is December 31, 1997. Forecasts of one-period ahead monthly variances are made from the parameter estimates by summing the individual one-period ahead through N_m-period ahead daily forecasts, where N_m represents the number of trading days in the relevant month. Numbers in parentheses under the parameter estimates represent standard errors. Parameters which are significantly different from zero at the 5% level are indicated with an asterisk.

Table 7.31: Calibration parameters (RD-GARCH, five year rolling)

Ending	γ_0	γ_2	γ_3	γ_4	λ	ν	b	ω	α	β
December 29, 1995	$3.641^{[-4]}$	0.101	0.002	$6.938^{[-4]}$	0.062	4.881	1.455	0.003	$1.122^{[-6]}$	0.995
	$(1.724^{[-4]})$*	(0.030)*	(0.031)	$(3.966^{[-4]})$	(0.396)	(1.260)*	(2.804)	(0.005)	$(8.313^{[-6]})$	(0.004)*
					0	0	0	0	0	0
					0	0	0	0	0	0
March 29, 1996	$2.625^{[-4]}$	0.121	0.019	$6.570^{[-4]}$	3.234	1.593	1.040	$7.409^{[-9]}$	0.189	0.582
	$(1.669^{[-4]})$	(0.034)*	(0.035)	$(3.965^{[-4]})$	(1.457)*	(0.781)*	(0.377)*	$(6.224^{[-8]})$	(0.195)	(0.178)*
					0	0	0	0	0	−0.582
										(0.003)*
					0	0	0	0	0	0
June 28, 1996	$2.331^{[-4]}$	0.139	0.017	$8.238^{[-4]}$	$1.310^{[-6]}$	1.175	1.151	0.257	$1.145^{[-7]}$	0.743
	$(1.682^{[-4]})$	(0.031)*	(0.035)	$(3.954^{[-4]})$*	$(4.087^{[-7]})$*	(0.391)*	(0.131)*	$(6.314^{[-7]})$*	$(8.976^{[-8]})$	$(6.668^{[-10]})$*
					0	0	0	$-1.295^{[-7]}$	0	0
								$(9.199^{[-20]})$*		
					0	0	0	0	0	0
September 30, 1996	$2.324^{[-4]}$	0.138	0.030	$7.820^{[-4]}$	3.514	1.798	1.148	$1.598^{[-9]}$	0.143	0.605
	$(1.690^{[-4]})$	(0.031)*	(0.033)	$(3.949^{[-4]})$	(1.181)*	(0.494)*	(0.382)*	$(1.081^{[-8]})$	(0.096)	(0.179)*
					0	0	0	0	0	−0.605
										(0.003)*
					0	0	0	0	0	0
December 31, 1996	$2.121^{[-4]}$	0.138	0.005	$6.868^{[-4]}$	3.646	2.659	2.208	$5.153^{[-10]}$	0.031	0.572
	$(1.665^{[-4]})$	(0.030)*	(0.031)	$(3.806^{[-4]})$	(1.638)*	(0.480)*	(0.810)*	$(4.445^{[-9]})$	(0.025)	(0.193)*
					0	0	0	0	0	0
					0	0	0	0	0	0

Entries of the form $x^{[-y]}$ represent scientific notation; that is $2.5^{[-2]}$ represents the value 2.5×10^{-2}.

Table 7.31: continued.

Ending	γ_0	γ_2	γ_3	γ_4	λ	ν	b	ω	α	β
March 31, 1997	$2.767^{[-4]}$	0.165	0.004	$5.302^{[-4]}$	3.914	2.614	1.940	$1.423^{[-10]}$	0.043	0.562
	$(1.676^{[-4]})$	(0.030)*	(0.031)	$(3.766^{[-4]})$	(0.903)*	(0.566)*	(0.927)*	$(7.021^{[-10]})$	(0.054)	(0.135)*
					0	0	0	0	0	−0.510
										(0.153)*
					0	0	0	0	0	0
June 30, 1997	$3.953^{[-4]}$	0.166	−0.010	$3.601^{[-4]}$	$1.379^{[-6]}$	1.474	0.727	0.063	$7.424^{[-8]}$	0.937
	$(1.679^{[-4]})$*	(0.031)*	(0.031)	$(3.731^{[-4]})$	$(5.593^{[-7]})$*	(0.259)*	(0.138)*	$(1.979^{[-7]})$*	$(2.221^{[-13]})$*	$(5.404^{[-22]})$*
					0	0	0	0	0	0
					0	0	0	0	0	0
September 30, 1997	$4.080^{[-4]}$	0.162	−0.006	$4.348^{[-4]}$	6.312	2.780	0.865	$2.743^{[-16]}$	0.157	0.634
	$(1.703^{[-4]})$*	(0.030)*	(0.032)	$(3.810^{[-4]})$	(0.064)*	(0.233)*	(0.092)*	$(8.410^{[-20]})$*	$(3.311^{[-6]})$*	$(4.313^{[-6]})$*
					0	0	0	0	0	−0.634
										$(5.459^{[-5]})$*
					0	0	0	0	0	0

Entries of the form $x^{[-y]}$ represent scientific notation; that is $2.5^{[-2]}$ represents the value 2.5×10^{-2}.

The table presents the parameter estimates which result when the RD-GARCH model is estimated for a five year rolling calibration period. In each case, the calibration period encompasses the five calendar years period which ends as indicated in the first column of the table. The model parameters are then used to forecast the one-period ahead conditional variance for the trading days which begin 1 day following the date indicated in the first column, and end on the date indicated in the first column of the next row of the table. The final date for which a one-period ahead daily forecast is made is December 31, 1997. Forecasts of one-period ahead monthly variances are made from the parameter estimates by summing the individual one-period ahead through N_m-period ahead daily forecasts, where N_m represents the number of trading days in the relevant month. Numbers in parentheses under the parameter estimates represent standard errors. Parameters which are significantly different from zero at the 5% level are indicated with an asterisk. The first pair of rows in each section presents the parameter estimates and standard errors for the base regime, and the remaining four rows present the parameter estimates and standard errors for regimes of unusually low and high volatility respectively. Parameters which are excluded by the stepwise estimation procedure are represented by a value of zero, with no corresponding standard error presented.

Chapter 8

Conclusions

This study develops the Regime Dependent Generalized Auto Regressive Conditional Heteroskedasticity (RD-GARCH) model and applies it to a daily index of returns on U.S. equities covering the period 1926 to 1997. The model is based on the Auto Regressive Conditional Heteroskedasticity (ARCH) methodology, first suggested by Engle (1982) and refined by subsequent researchers. The ARCH approach provides models comprising two parts: a conditional mean equation which defines the evolution of a dependent variable in terms of one or more independent variables; and a conditional variance equation which defines the evolution of the time-varying variance of the residuals from the conditional mean equation in terms of current and prior unexpected shocks, prior conditional volatility, and possibly other exogenous variables. The RD-GARCH model is different from previous models in that it combines the ARCH methodology with a general approach for allowing model parameters to vary across regimes of differing volatility. In addition to the development of the model, a number of additional issues related to the estimation and application of the model have been addressed.

GARCH models are commonly estimated by constructing a likelihood function for the model and maximizing that likelihood function using numerical optimization techniques. The maximum likelihood estimation technique requires the specification of a distributional assumption for the residuals from the conditional mean equation. Initial development of the model was performed under the assumption of normality.

One problem with estimating non-linear models is that many traditional optimization techniques tend to perform poorly in that they converge to a local extremum of the objective function rather than to the global maximum or minimum. In Chapter 2 a hybrid optimization methodology was developed which combined a robust, but slow, stochastic global optimization technique with a much faster deterministic local optimization technique. This hybrid technique was shown to provide an improvement, as measured by larger values for the likelihood function, over the conventional approach of starting the deterministic technique at a "plausible" starting point. The improvement did come at the expense of a ten-fold increase in the computer time taken for optimization, although the stochastic technique alone is considerably more expensive in computer time.

Because it nests many of the more popular GARCH models into a single equation, the Nested GARCH model of Hentschel (1995) was chosen as the foundation upon which to construct the RD-GARCH model. The Nested GARCH specification allows for a number of different effects in the conditional volatility, including two different approaches to modeling asymmetry in the response of conditional volatility to unexpected shocks of different signs. These two alternative approaches are characterized by a shift or a rotation of the news impact curve (NIC), which is a graphical device that illustrates the effect of the prior period unexpected shock on the current conditional volatility. The results from Chapter 3 indicated that the model which allowed the NIC to shift but where parameter restrictions were made that did not allow for rotation was an appropriate starting point for development of the RD-GARCH model.

The choice of model was reinforced by additional tests performed on the residuals from the conditional mean equation. The residuals were first standardized to provide zero mean and unit variance. The mean and variance of the standardized residuals were not significantly different from zero and one respectively, although they were significantly leptokurtic and negatively skewed. A Jarque-Bera test was performed which rejected the null hypothesis of normality. In addition, sign bias tests indicated that both the sign of the the previous shock and the magnitude of the previous negative shock retain some predictive power for current conditional volatility, although the magnitude of the previous non-negative shock did not have predictive power. These results indicated either misspecification of the model, or that some alternative distributional assumption for the disturbances would be appropriate.

Enhancement of the model was addressed in two ways. The mean equation was considered first, with the simple GARCH-M specification being replaced by a specification for the conditional mean which includes the return lagged one and two periods and a Monday dummy. The GARCH-M term was found to contribute no additional value in the presence of the other terms, and was excluded from the final model.

The conditional volatility equation was then addressed, as investigation of the standardized residuals from the mean equation indicated that, even with the best GARCH equation from the range of nested models, heteroskedasticity remained. The approach employed to address the remaining heteroskedasticity allowed the parameters of the GARCH equation to change during periods of volatility which were either unusually large or unusually small. This provision for time-varying GARCH equation parameters represents the enhancement to the GARCH methodology provided by the RD-GARCH model. The RD-GARCH model combines the evolutionary approach to modeling conditional volatility provided by the GARCH methodology with a provision for the unconditional volatility of the return series to undergo abrupt changes at discrete points in time.

Unfortunately, the rich parameterization of the variance equation provided by the regime dependent framework led to a problem of correlation among the explanatory variables. To address this issue, a simple stepwise procedure, analogous to the backward elimination method for the multiple linear regression model, was employed. This approach allows for the parsimonious selection of the appropriate parameters for a particular modeling situation, without the computational overhead required by estimation of all possible subsets of model parameters.

Both the RD-GARCH methodology itself and the particular RD-GARCH model identified as the most appropriate were then subject to a number of additional diagnostic tests to determine how well the approach performs at describing the volatility of the U.S. equity index over a particular time period.

The regime detection procedure was investigated first, providing the anecdotal observation that periods of high volatility in the equity index have, in general, been associated with depressions and recessions in the economy, and that periods of low volatility have been associated with the involvement of the United States in wars.

The RD-GARCH methodology was applied to the full set of 20 models provided by the Nested GARCH specification. In all cases, the regime dependent model was found to be superior to the conventional model. The RD-GARCH methodology resulted in a model fit which was superior to that obtained by an otherwise identical model which did not allow the GARCH parameters to vary across volatility regimes.

The particular RD-GARCH model identified as the most appropriate was then investigated in more detail. The conditional mean equation parameters were similar for both the conventional version of the model and the version of the model where the conditional volatility equation parameters were allowed to vary across regimes of differing volatility. In both cases, the impact of the prior period's return was positive, the return two periods ago had a small negative impact, and the Monday effect was very small and negative. The conditional variance parameters showed considerably less similarity between the conventional and regime dependent versions of the model. The parameter λ, which controls the Box-Cox transformation of the conditional volatility series and has the effect of transforming the functional form of the conditional volatility, had a value of approximately one for the conventional version of the model. This indicated that no alternative functional form provided by the Box-Cox transformation was effective in improving the normality or homoskedasticity of the conditional volatility. When λ was allowed to vary across regimes, it took on a value of nearly five in the base regime, but a value of approximately one again for both the high and low volatility regimes. Either the separation of the series into regimes removed the problems of non-normality and heteroskedasticity for the periods of unusually high or low volatility, or the nature of the non-normality and heteroskedasticity was different in these regimes and the Box-Cox transformation was not effective in rendering the series more nearly homoskedastic and normally distributed.

The NIC for the model in which regime detection and isolation was not performed was shifted somewhat to the right, and exhibited a small amount of convexity; positive unexpected shocks had a greater impact on subsequent conditional volatility than negative shocks of the same magnitude, and the response of the conditional volatility to volatility shocks of the same sign but different magnitudes was not linear. The difference in impact between a shock with a magnitude of two standard deviations and a shock of three standard deviations was greater than the difference in impact between a shock of one standard deviation and a shock of two standard deviations. If the volatility shocks are viewed as a response to news, these results indicate that a positive news event had a smaller effect on subsequent conditional volatility than a negative news event of the same magnitude, and larger news events of either sign have a disproportionately larger impact on subsequent conditional volatility.

The NIC for the regime dependent version of the model comprises three separate curves. The convexity exhibited by the NIC differed markedly across the three different regime types. High volatility regimes were characterized by an NIC with a much higher degree of curvature than the curve for either the base regime or low volatility regimes. Conditional volatility is more responsive to large shocks in periods of high volatility, and this effect is not simply due to the higher unconditional volatility that exists in these regimes. Low volatility regimes exhibited a slightly sharper NIC than the base regime, although the NIC exhibited less curvature. For a certain range in magnitude of shocks, the difference in impact for slightly larger shocks is greater in low volatility periods than in normal volatility periods. Once outside that range, however, the reverse is true; the difference in impact for much larger shocks is greater in normal volatility periods than in low volatility periods. The curvature of the NIC for the conventional version of the model appeared to be an average of the curvature of the three NICs from the regime dependent model. Allowing the model parameters to vary across regimes provided a mechanism for identifying a difference in the way the market reacts to news during periods of differing overall volatility. The degree to which the NIC is shifted to the right was similar for all three curves in the regime dependent model. The degree of asymmetry was similar across regime types.

Two results of interest obtained from additional tests performed on the residuals from the model: First, introduction of regime dependence resulted in a notable decrease in the level of excess kurtosis observed in the standardized residuals, although the residuals did continue to exhibit statistically significant leptokurtosis. Secondly, the sign bias tests indicated that the augmented mean equation[1] alone was able to remove the predictive power of the sign of the previous shock, but was unable to account completely for the predictive power of the magnitude of the previous negative shock. Modeling regime dependence in the conditional volatility with the RD-GARCH model, however, was able to remove the predictive power of the magnitude of the previous negative shock.

In an effort to address the residual leptokurtosis issue, two alternative distributional assumptions were examined. Likelihood functions constructed for the generalized error distribution and Student-*t* distribution both resulted in significantly greater likelihood values than for the likelihood function constructed for the normal distribution. Because the alternative models were not all constructed with parameter restrictions on a single model, likelihood ratio tests were unable to provide an indication as to which of the two alternative distributions would be preferred for the conventional version of the models, and could not provide any information regarding the regime dependent models. The sign bias tests indicated that an assumption of normality for disturbances was appropriate in the case of regime dependent models, and did not suggest that one distribution was preferred over any other for the conventional versions of the model. In light of this ambiguity regarding the appropriate distributional assumption, the ongoing analysis of the RD-GARCH methodology was conducted under the assumption of normality for the residuals.

---

[1] The conditional mean equation which included the return lagged one and two periods, a Monday dummy, but not the GARCH-M term.

Several issues related to the stability of the RD-GARCH methodology were then investigated. The question of excess vs. raw returns was the first issue addressed, both because of the availability of a much longer raw returns series than the available excess returns series and because the construction of the excess return series required the extrapolation of a daily risk-free rate from a monthly rate. As there was not a notable difference between the results obtained with the raw return series and those obtained with the excess return series, and as the raw return series avoided the problems associated with using a monthly risk-free series to construct a daily excess return series, the raw return series was adopted.

The entire sample period was extended to comprise the dates January 2, 1915 through December 29, 1995. This extended sample was then broken into two subsamples, and each of the two subsamples further divided into two additional sub-periods. The seven different sample periods were then investigated both from the point of view of the regime detection and isolation procedure and from the point of view of the stability of the estimated model parameters.

In general, with the exception of the 1929–1951 sub-period, the regime detection procedure adopted for this study was stable and located essentially the same regimes regardless of the extent of the sample period. As the sample period which began at the beginning of the fourth quarter of 1929 and ended at the end of 1951 was dominated by the unique and extremely high volatility observed during the Great Depression, and as its occurrence was sufficiently far in the past that the domination of the regime selection process was likely not to be an issue for any forecasting applications, the sensitivity of the process to the high volatility regime was judged to be not material.

Examination of the NICs and HICs for the conventional version of the model indicated that, with the exception once again of the extreme volatility period associated with the 1930's, the nature of the impact of news on subsequent conditional volatility was reasonably stable across sample periods. The results from the application of the RD-GARCH model, however, were interesting in their diversity. In particular, the curves from the regime dependent model applied to the four small sub-periods were all notably different from each other. This observation suggests that fitting a single GARCH model with time-invariant parameters to a long time series may well be inappropriate. Allowing the parameters to vary across regimes provided additional insight into the nature of the volatility in the market.

As a final test of the RD-GARCH methodology, several alternative models were used to construct daily one-period ahead and monthly one-period ahead forecasts of volatility for a hold-out sample comprising the period January 2, 1996 to December 31, 1997. In addition to testing several different models of volatility, three different calibration periods were used. A number of tests were performed which compared the results of the alternative forecasting models to those obtained with the RD-GARCH model.

The results for the daily forecasts were similar for all three calibration periods. In all cases, the RD-GARCH model provided the best fit under the mean square error (MSE) criterion, and in two of the three cases it also provided the best fit under the mean absolute error (MAE) criterion. In the third case, the RD-GARCH model was second best under the MAE criterion. The RD-GARCH methodology also appeared to be insensitive to the length of the calibration period, provided the period was sufficiently long that regime detection was effective. The tests of sufficiency of RD-GARCH and tests for model improvement by RD-GARCH suggested that forecasts of daily variances could be improved by the inclusion of a random walk component in the forecast model. The tests for difference in MSE suggested that, with a single exception where the RD-GARCH model was superior to the alternative model, there was not a statistically significant difference in MSE among the forecast models. In general, however, the RD-GARCH model provided lower forecast error variance than the alternative models at the expense of greater forecast bias. As it is likely that low forecast error variance is preferable to low forecast bias, and as the longer calibration period performed slightly better than the shorter periods with very little extra cost, the RD-GARCH forecast model calibrated with the longest calibration period was judged to be the preferred model for forecasting daily return variance.

The results for monthly forecasts exhibited considerably less consistency across the calibration periods than was shown by the results for the daily forecasts. The minimum MSE and MAE for the longest calibration period were again provided by the RD-GARCH methodology, with different models preferred for both shorter periods. Overall, the RD-GARCH model provided the lowest MAE for two of the three calibration periods, and the lowest MSE for the longest calibration period. The preferred model in the remaining cases was different for each calibration period. The RD-GARCH model calibrated with the longest calibration period also provided the lowest overall MAE and the second lowest overall MSE.

The tests of sufficiency of RD-GARCH and tests for model improvement by RD-GARCH suggested that for both longer calibration periods the forecasts of monthly variance provided by the RD-GARCH model would likely not be improved by a model combining the RD-GARCH methodology with any of the other models. Although a composite model combining a medium length moving average component with the RD-GARCH model was suggested for the shortest calibration period, the longer calibration periods performed better overall than the shortest calibration period, indicating that it would be more appropriate to use the longer calibration period rather than attempting to construct a such a composite model. The tests for difference in MSE indicated that in ten of the 21 cases the RD-GARCH provided an MSE which was significantly less than that provided by the alternative model. There were two cases where the reverse was true, however. In general the RD-GARCH model provided lower forecast error variance than the alternative models at the expense of greater forecast bias. The RD-GARCH model also provided the greatest consistency in the forecasting of monthly return variance. It was the preferred model for the longest calibration period, and the second best model for both of the shorter calibration periods. The preferred model was different for the each of the shorter updating calibration period. In light of this consistency, and because of its overall performance, the RD-GARCH forecast model calibrated with the longest calibration period was also judged to be the preferred model for forecasting monthly return variance.

As a final test of forecast performance, RD-GARCH models were constructed under two alternative distributions for the disturbances: the generalized error distribution and the Student-t distribution. This test addressed the lack of a clear preference for one of the three alternative distributions noted earlier for within-sample fit. For the daily one-period ahead forecasts, the assumption of disturbances from the Student-t distribution provided forecasts which were superior to those provided under the assumption both of normally distributed disturbances and of disturbances from the generalized error distribution. For forecasts of monthly return variance, the model constructed under the assumption of normally distributed errors provided the superior forecast.

The Regime Dependent Auto Regressive Conditional Heteroskedasticity model developed in this study has been shown to provide better within-sample fit than a large group of alternative models from a the GARCH family. The methodology also provides a mechanism for comparing and interpreting the relationship between unexpected shocks and subsequent conditional volatility during periods of differing overall market volatility. The model performs very well at out-of-sample forecasting, and is insensitive to the length of the period used for calibration. Taken together, these results indicate that combining a model which allows volatility to change in an evolutionary manner with a model which allows for abrupt changes in volatility at discrete points in time provides a significant improvement over previously available models. The RD-GARCH model provides a general approach to combining these two alternative modeling methodologies, thereby providing a valuable addition to the GARCH family of conditional volatility models.

Bibliography

Vedat Akgiray. Conditional heteroskedasticity in time series of stock returns: Evidence and forecasts. *Journal of Business*, 62(1):55–80, 1989.

R. Ashley, C.W. Granger, and R. Schmalensee. Advertising and aggregate consumption: An analysis of causality. *Econometrica*, 48(5):1149–1167, July 1980.

Brad Baldauf and G. J. Santoni. Stock price volatility : Some evidence from an ARCH model. *Journal of Futures Markets*, 11(2):191–200, April 1991.

E. K. Berndt, B. H. Hall, R. E. Hall, and J. A. Hausman. Estimation and inference in nonlinear structural models. *Annals of Economic and Social Measurement*, 3: 653–665, 1974.

Steve Beveridge. Time series analysis and forecasting. University of Alberta Managment Science 614 lecture notes, 1994.

Christian Bischof, Alan Carle, Peyvand Khademi, Andrew Mauer, and Paul Hovland. Adifor 2.0 user's guide. Technical Memorandum 192, Mathematics and Computer Science Division, Argonne National Laboratory, 9700 S. Cass Ave., Argonne, IL 60439, August 1995.

T. Bollerslev, R.Y. Chou, and K.F. Kroner. ARCH modeling in finance: A review of the theory and empirical evidence. *Journal of Econometrics*, 52:5–59, 1992.

Tim Bollerslev. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3):307–327, April 1986.

Tim Bollerslev and Jeffrey M. Woolridge. Quasi-maximum likelihood estimation and inference in dynamic models with time-varying covariances. *Econometric Reviews*, 11:143–172, 1992.

George E. P. Box and Gwilym M. Jenkins. *Time series analysis: forecasting and control.* Holden-Day, San Francisco, revised edition, 1976.

Chris Brooks. Predicting stock index volatility: Can market volume help? *Journal of Forecasting*, 17:59–80, 1998.

Adolph Buse. Econometrics I. University of Alberta Economics 507 lecture notes, 1992.

Adolph Buse. Econometrics II. University of Alberta Economics 508 lecture notes, 1993.

John Y. Campbell and Ludger Hentschel. No news is good news : An asymmetric model of changing volatility in stock returns. *Journal of Financial Economics*, 31 (3):281–318, June 1992.

John Y. Campbell, Andrew W. Lo, and A. Craig MacKinlay. *The Econometrics of Financial Markets.* Princeton University Press, Princeton, New Jersey, 1997.

Robert Dacco and Steve Satchell. Why do regime-switching models forecast so badly? *Journal of Forecasting*, 18:1–16, 1999.

Theodore E. Day and Craig M. Lewis. Stock market volatility and the information content of stock index options. *Journal of Econometrics*, 52:267–287, 1992.

Zhuanxin Ding, Clive W. J. Granger, and Robert F. Engle. A long memory property of stock market returns and a new model. *Journal of Empirical Finance*, 1:83–106, 1993.

R. Engle and M. Rothschild. ARCH models in finance. *Journal of Econometrics*, 52: 245–266, 1992.

Robert F. Engle. Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 50:987–1007, 1982.

Robert F. Engle, David M. Lilien, and Russell P. Robins. Estimating time varying risk premia in the term structure: The ARCH-M model. *Econometrica*, 55(2): 391–407, March 1987.

Robert F. Engle and Victor K. Ng. Measuring and testing the impact of news on volatility. *Journal of Finance*, 48(5):1749–1778, December 1993.

Estima. Rats version 4. Computer Software, 1995.

Philip Hans Franses and Dick Van Dijk. Forecasting stock market volatility using (non-linear) GARCH models. *Journal of Forecasting*, 15:229–235, 1996.

Kenneth R. French. Stock returns and the weekend effect. *Journal of Financial Economics*, 8:55–69, 1980.

Kenneth R. French, G. William Schwert, and Robert F. Stambaugh. Expected stock returns and volatility. *Journal of Financial Economics*, 19:3–29, 1987.

John E. Freund. *Mathematical Statistics.* Prentice-Hall, Englewood Cliffs, N.J., 1962.

David M. Gay. Algorithm 611—subroutines for unconstrained minimization using a model/trust-region approach. *ACM Transactions on Mathematical Software*, 9: 503–524, 1983.

Lawrence R. Glosten, Ravi Jagannathan, and David E. Runkle. On the relation between the expected value and the volatility of the nominal excess return on stocks. *Journal of Finance*, 48(5):1779–1801, December 1993.

Willian L. Goffe, Gary D. Ferrier, and John Rogers. Global optimization of statistical functions with simulated annealing. *Journal of Econometrics*, 60:65–99, 1994.

Stephen F. Gray. Modeling the conditional distribution of interest rates as a regime-switching process. *Journal of Financial Economics*, 42(1):27–62, September 1996.

James D. Hamilton. A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57(2):357–84, March 1989.

James D. Hamilton and Raul Susmel. Autoregressive conditional heteroskedasticity and changes in regime. *Journal of Econometrics*, 64(1):307–333, September/October 1994.

Robert A. Haugen, Eli Talmor, and Walter N. Torous. The effect of volatility changes on the level of stock prices and subsequent expected returns. *Journal of Finance*, 46(3):985–1007, July 1991.

Ludger Hentschel. All in the family: Nesting symmetric and asymmetric GARCH models. *Journal of Financial Economics*, 39(1):71–104, September 1995.

Matthew L. Higgins and Anil K. Bera. A class of nonlinear ARCH models. *International-Economic-Review*, 33(1):137–58, February 1992.

Carlos M. Jarque and Anil K. Bera. A test for normality of observations and regression residuals. *International Statistical Review*, 55(2):163–172, 1987.

J.D. Jobson. *Applied Multivariate Data Analysis*, volume I: Regression and Experimental Design. Springer-Verlag, New York, NY, 1991.

George G. Judge, W. E. Griffiths, R. Carter Hill, Helmut Lütkepohl, and Tsoung-Chao Lee. *The Theory and Practice of Econometrics*. Wiley, New York, NY, second edition, 1985.

Christopher G. Lamoureaux and William D. Lastrapes. Forecasting stock-return variance: Towards an understanding of stochastic implied volatilities. *The Review of Financial Studies*, 6(2):293–326, 1993.

H. Levene. Robust tests for equality of variances. In Ingram Olkin, editor, *Contributions to Probability and Statistics*, pages 278–292. Stanford University Press, Palo Alto, California, 1960.

Andrew W. Lo and A. Craig MacKinlay. An econometric analysis of nonsynchronous trading. *Journal of Econometrics*, 45:181–211, 1990.

J. A. Nelder and R. Mead. A simplex method for function minimization. *The Computer Journal*, 7:308–313, 1965.

Daniel B. Nelson. Conditional heteroskedasticity in asset returns : A new approach. *Econometrica*, 59(2):347–370, March 1991.

W. Newey and K. West. A simple positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55:703–708, 1987.

Editor Phyllis A. Fox. *The PORT Mathematical Subroutine Library.* AT&T Bell Laboratories, Inc., Murray Hill, NJ, 3.0 edition, May 1984.

Kenneth Price and Ranier Storn. Differential evolution: A simple evolution strategy for fast optimization. *Dr. Dobb's Journal*, pages 18–24, April 1997.

G. William Schwert. Indexes of united states stock prices from 1802 to 1987. *Journal of Business*, 63:399–426, July 1989.

Ruey S. Tsay. Outliers, level shifts, and variance changes in time series. *Journal of Forecasting*, 7:1–20, 1988.

Waterloo Maple Inc. Maple V Release 5. Computer software, 1998.

Robert F. Whitelaw. Time variations and covariations in the expectation and volatility of stock market returns. *Journal of Finance*, 49(2):515–541, June 1994.

Appendix A

Model Summary

Table A.1: Summary of model definitions

Model Name	Conditional Mean	Conditional Volatility Restrictions	Distribution	Common Name
model01	$r_t = \gamma_0 + \gamma_1 h_t^2 + u_t$	$\lambda = 0, \nu = 1, b = 0, c = 0$	Normal	—
model02	$r_t = \gamma_0 + \gamma_1 h_t^2 + u_t$	$\lambda = 0, \nu = 1, c = 0$	Normal	—
model03	$r_t = \gamma_0 + \gamma_1 h_t^2 + u_t$	$\lambda = 0, \nu = 1, b = 0$	Normal	Related to EGARCH[1]
model04	$r_t = \gamma_0 + \gamma_1 h_t^2 + u_t$	$\lambda = 0, \nu = 1$	Normal	—
model05	$r_t = \gamma_0 + \gamma_1 h_t^2 + u_t$	$\lambda = 1, \nu = 1, b = 0, c = 0$	Normal	—
model06	$r_t = \gamma_0 + \gamma_1 h_t^2 + u_t$	$\lambda = 1, \nu = 1, c = 0$	Normal	—
model07	$r_t = \gamma_0 + \gamma_1 h_t^2 + u_t$	$\lambda = 1, \nu = 1, b = 0$	Normal	—
model08	$r_t = \gamma_0 + \gamma_1 h_t^2 + u_t$	$\lambda = 1, \nu = 1$	Normal	AGARCH
model09	$r_t = \gamma_0 + \gamma_1 h_t^2 + u_t$	$\lambda = 2, \nu = 2, b = 0, c = 0$	Normal	GARCH
model10	$r_t = \gamma_0 + \gamma_1 h_t^2 + u_t$	$\lambda = 2, \nu = 2, c = 0$	Normal	N-A GARCH
model11	$r_t = \gamma_0 + \gamma_1 h_t^2 + u_t$	$\lambda = 2, \nu = 2, b = 0$	Normal	GJR GARCH
model12	$r_t = \gamma_0 + \gamma_1 h_t^2 + u_t$	$\lambda = 2, \nu = 2$	Normal	—
model13	$r_t = \gamma_0 + \gamma_1 h_t^2 + u_t$	$\lambda = \nu, b = 0, c = 0$	Normal	NARCH
model14	$r_t = \gamma_0 + \gamma_1 h_t^2 + u_t$	$\lambda = \nu, c = 0$	Normal	—
model15	$r_t = \gamma_0 + \gamma_1 h_t^2 + u_t$	$\lambda = \nu, b = 0$	Normal	APARCH
model16	$r_t = \gamma_0 + \gamma_1 h_t^2 + u_t$	$\lambda = \nu$	Normal	—
model17	$r_t = \gamma_0 + \gamma_1 h_t^2 + u_t$	$b = 0, c = 0$	Normal	—
model18	$r_t = \gamma_0 + \gamma_1 h_t^2 + u_t$	$c = 0$	Normal	Extended GARCH
model19	$r_t = \gamma_0 + \gamma_1 h_t^2 + u_t$	$b = 0$	Normal	—
model20	$r_t = \gamma_0 + \gamma_1 h_t^2 + u_t$	—	Normal	—

Table A.1: continued.

Model Name	Conditional Mean	Conditional Volatility Restrictions	Distribution	Common Name
model21	$r_t = \gamma_0 + \gamma_2 r_{t-1} + \gamma_3 r_{t-2} + \gamma_4 D_M + u_t$	$b = 0, c = 0$	Normal	—
model22	$r_t = \gamma_0 + \gamma_2 r_{t-1} + \gamma_3 r_{t-2} + \gamma_4 D_M + u_t$	$c = 0$	Normal	Extended GARCH
model23	$r_t = \gamma_0 + \gamma_2 r_{t-1} + \gamma_3 r_{t-2} + \gamma_4 D_M + u_t$	$b = 0$	Normal	—
model24	$r_t = \gamma_0 + \gamma_2 r_{t-1} + \gamma_3 r_{t-2} + \gamma_4 D_M + u_t$	—	Normal	—
model25	$r_t = \gamma_0 + \gamma_1 h_t^2 + \gamma_2 r_{t-1} + \gamma_3 r_{t-2} + \gamma_4 D_M + u_t$	$c = 0$	Normal	Extended GARCH
model26	$r_t = \gamma_0 + \gamma_2 r_{t-1} + \gamma_4 D_M + u_t$	$c = 0$	Normal	Extended GARCH
model27	$r_t = \gamma_0 + \gamma_2 r_{t-1} + u_t$	$c = 0$	Normal	Extended GARCH
model28	$r_t = \gamma_0 + \gamma_2 r_{t-1} + \gamma_3 r_{t-2} + u_t$	$c = 0$	Normal	Extended GARCH
model29	$r_t = \gamma_0 + \gamma_2 r_{t-1} + \gamma_3 r_{t-2} + \gamma_4 D_M + u_t$	$c = 0$	GED	Extended GARCH
model31	$r_t = \gamma_0 + \gamma_2 r_{t-1} + \gamma_3 r_{t-2} + \gamma_4 D_M + u_t$	$c = 0$	Student-t	Extended GARCH
model34	$r_t = \gamma_0 + \gamma_2 r_{t-1} + \gamma_3 r_{t-2} + \gamma_4 D_M + u_t$	$\lambda = 2, \nu = 2, b = 0, c = 0$	Normal	GARCH

The table presents a summary relating the model names used in the text and tables of this work to the conditional mean equations and conditional volatility equation parameter restrictions which define the various models. The fourth column presents the distributional assumption for disturbances under which the model was developed, and the fifth column presents the common name for the model if such a name exists. The models named "model30", "model32", and "model33" are not presented in the table as these models were used in earlier drafts of this work, and do not appear in the final version.

[1] model03 results in the conditional volatility equation $\ln h_t^2 = \omega + \alpha[|v_t| - cv_t] + \beta \ln h_{t-1}^2$. The EGARCH model of Nelson (1991), however, provides the conditional volatility equation $\ln h_t^2 = \omega + \alpha[|v_t| - E|v_t| + cv_t] + \beta \ln h_{t-1}^2$. Computing an estimate of the term $E|v_t|$ requires a distributional assumption about the disturbances, however, so the adjustment has not been made in order to maintain generality. In this sense, model03 is related to but not identical to Nelson's EGARCH.

Appendix B

Regime Detection Example

This appendix provides a simple example of the regime detection methodology. Consider the (hypothetical) return series:

r_1	r_2	r_3	r_4	r_5	r_6	r_7	r_8	r_9	r_{10}
−0.02	0.03	−0.04	0.01	0.02	0.03	−0.02	−0.04	0.03	−0.01

working with a window size of 2, results in the following variance ratios:

Window	1, 2	2, 3	3, 4	4, 5	5, 6	6, 7	7, 8	8, 9	9, 10
$\hat{R}_d$	1.47	3.57	1.58	0.05	0.06	1.47	0.28	3.570	0.91
$\hat{R}_d^{-1}$	0.68	0.28	0.63	18.86	15.86	0.68	3.56	0.28	1.10

The window which includes r_4 and r_5 gives the largest value for either $\hat{R}_d$ or $\hat{R}_d^{-1}$ of 18.86, so this is the window which is selected for further investigation. Adding one observation onto the beginning and end of this window independently yields:

Window	3, 4, 5	4, 5, 6
$\hat{R}_d$	1.19	0.84
$\hat{R}_d^{-1}$	0.10	10.19

Adding a single observation to the end of the window made the smallest change to $\hat{R}_d^{-1}$ (18.86 to 10.19) so this observation is kept in the window, and the process is repeated:

Window	3, 4, 5, 6	4, 5, 6, 7
$\hat{R}_d$	2.10	0.48
$\hat{R}_d^{-1}$	0.83	1.21

This time, adding a single observation to either end of the window results in a variance ratio which is smaller than the critical value of 3.0, so the process is stopped and a *low* regime is reported covering observations 4–6. The procedure would then be repeated, excluding any potential regimes which contained an observation which had already been included in another high or low regime, until no more windows resulted in a value for $\hat{R}_d$ or $\hat{R}_d^{-1}$ larger than the critical value.

Appendix C

Log-Likelihood Derivation

C.1 Standard Normal Distribution

For the Standard Normal distribution, where v_t has zero mean and unit variance, the density of v_t is

$$f(v_t) = -\frac{1}{\sqrt{2\pi}} \exp\left[-\frac{1}{2}v_2^2\right] \tag{C.1}$$

Using the formulation for a change in variables:

$$\begin{array}{rl} \text{if} & x_1 \sim f_1(x_1) \\ \text{and} & x_1 = h(x_2) \\ \text{then} & x_2 \sim f_1(h(x_2)) \left|\dfrac{\partial h(x_2)}{\partial x_2}\right| \end{array} \tag{C.2}$$

The density of r_t is given by:

$$\begin{aligned} f(r_t) &= -\frac{1}{\sqrt{2\pi}} \exp\left[-\frac{1}{2}F_t(r_t,\gamma)^2\right] \frac{\partial F_t(r_t,\gamma)}{\partial r_t} \\ &= -\frac{1}{\sqrt{2\pi}} \exp\left[-\frac{v_t^2}{2}\right] \frac{1}{h_t} \end{aligned} \tag{C.3}$$

The log-likelihood is:

$$\mathcal{L} = \sum_{t=1}^{T} \left[-\frac{1}{2}\log 2\pi - \frac{1}{2}v_t^2 - \log h_t\right] \tag{C.4}$$

C.2 Generalized Error Distribution

For the generalized error distribution, the density of v_t, where v_t has zero mean and unit variance, is:

$$f(v_t) = \frac{k \exp\left[-\frac{1}{2}\left|\frac{v_t}{\Lambda}\right|^k\right]}{\Lambda 2^{(1+\frac{1}{k})}\Gamma[1/k]} \tag{C.5}$$

where

$$\Lambda \equiv \left(2^{-\frac{2}{k}}\frac{\Gamma[1/k]}{\Gamma[3/k]}\right)^{\frac{1}{2}} \tag{C.6}$$

and k is a positive parameter which controls the shape of the tails. For the case where $k = 2$, the generalized error distribution reduces to the Standard Normal distribution. Applying the change of variables formulation from Equation C.2, the density of r_t is:

$$\begin{aligned} f(r_t) &= \frac{k \exp\left[-\frac{1}{2}\left|\frac{F_t(r_t,\gamma)}{\Lambda}\right|^k\right]}{\Lambda 2^{(1+\frac{1}{k})}\Gamma[1/k]}\frac{\partial F_t(r_t,\gamma)}{\partial r_t} \\ &= \frac{k \exp\left[-\frac{1}{2}\left|\frac{v_t}{\Lambda}\right|^k\right]}{\Lambda 2^{(1+\frac{1}{k})}\Gamma[1/k]}\frac{1}{h_t} \end{aligned} \tag{C.7}$$

The log-likelihood is:

$$\begin{aligned} \mathcal{L} = \sum_{t=1}^{T}\Bigg[&-\log 2 - \left|v_t\left(\frac{\Gamma[3/k]}{\Gamma[1/k]}\right)^{\frac{1}{2}}\right|^k \\ &+ \log k + \frac{1}{2}\log\Gamma[3/k] - \frac{3}{2}\log\Gamma[1/k] - \log h_t\Bigg] \end{aligned} \tag{C.8}$$

C.3 Student-*t* Distribution

For the standardized Student-t distribution, where v_t has zero mean and unit variance, the density of v_t is:

$$f(v_t) = \frac{\Gamma[(k+1)/2]}{\sqrt{\pi(k-2)}\Gamma[k/2]}\left(1+\frac{v_t^2}{k-2}\right)^{\frac{-(k+1)}{2}} \tag{C.9}$$

Where k represents the number of degrees of freedom. As $k \to \infty$, the density of the Student-t distribution approaches that of the Standard Normal distribution. For Applying the change of variables formulation from Equation C.2:

$$\begin{aligned} f(r_t) &= \frac{\Gamma[(k+1)/2]}{\sqrt{\pi(k-2)}\Gamma[k/2]}\left(1+\frac{F_t(r_t,\gamma)^2}{k-2}\right)^{\frac{-(k+1)}{2}} \frac{\partial F_t(r_t,\gamma)}{\partial r_t} \\ &= \frac{\Gamma[(k+1)/2]}{\sqrt{\pi(k-2)}\Gamma[k/2]}\left(1+\frac{v_t^2}{k-2}\right)^{\frac{-(k+1)}{2}} \frac{1}{h_t} \end{aligned} \tag{C.10}$$

The log-likelihood is:

$$\begin{aligned} \mathcal{L} = \sum_{t=1}^{T} \Bigg[-\frac{1}{2}\log \pi + \log \Gamma[(k+1)/2] - \log \Gamma[k/2] \\ +\frac{1}{2} k \log(k-2) - \frac{k+1}{2} \log\left(k-2+v_t^2\right) - \log h_t \Bigg] \end{aligned} \tag{C.11}$$

Appendix D

Supplementary Tables

Table D.1: Estimates of model parameters (alternative mean equations)

Model	γ_0	γ_1	γ_2	γ_3	γ_4	λ	ν	b	c	ω	α	β	Likelihood
model25	$5.857^{[-4]}$ $(7.355^{[-5]})^*$	-0.003 (0.903)	0.165 $(0.009)^*$	-0.035 $(0.008)^*$	-0.002 $(1.436^{[-4]})^*$	0.941 $(0.269)^*$	1.537 $(0.117)^*$	0.489 $(0.057)^*$	0	$1.867^{[-4]}$ $(2.434^{[-4]})$	0.059 $(0.017)^*$	0.926 $(0.018)^*$	59093.055
model18 (GARCH-M)	$3.863^{[-4]}$ $(8.377^{[-5]})^*$	-0.654 (0.977)	0	0	0	0.986 $(0.292)^*$	1.509 $(0.118)^*$	0.452 $(0.054)^*$	0	$1.558^{[-4]}$ $(2.189^{[-4]})$	0.064 $(0.019)^*$	0.924 $(0.020)^*$	58810.909
model27 (r_{t-1})	$2.228^{[-4]}$ $(5.601^{[-5]})^*$	0	0.154 $(0.009)^*$	0	0	1.030 $(0.290)^*$	1.573 $(0.123)^*$	0.526 $(0.061)^*$	0	$1.291^{[-4]}$ $(1.797^{[-4]})$	0.064 $(0.017)^*$	0.918 $(0.021)^*$	58986.964
model28 (r_{t-1}, r_{t-2})	$2.521^{[-4]}$ $(5.636^{[-5]})^*$	0	0.159 $(0.009)^*$	-0.035 $(0.009)^*$	0	1.005 $(0.285)^*$	1.553 $(0.121)^*$	0.505 $(0.060)^*$	0	$1.451^{[-4]}$ $(1.982^{[-4]})$	0.063 $(0.017)^*$	0.920 $(0.020)^*$	58996.327
model26 (r_{t-1}, D_M)	$5.578^{[-4]}$ $(5.720^{[-5]})^*$	0	0.160 $(0.009)^*$	0	-0.002 $(1.455^{[-4]})^*$	0.963 $(0.266)^*$	1.555 $(0.114)^*$	0.511 $(0.058)^*$	0	$1.686^{[-4]}$ $(2.150^{[-4]})$	0.060 $(0.016)^*$	0.924 $(0.019)^*$	59083.569
model22 (r_{t-1}, r_{t-2}, D_M)	$5.856^{[-4]}$ $(5.780^{[-5]})^*$	0	0.165 $(0.009)^*$	-0.035 $(0.008)^*$	-0.002 $(1.437^{[-4]})^*$	0.941 $(0.265)^*$	1.538 $(0.114)^*$	0.489 $(0.056)^*$	0	$1.861^{[-4]}$ $(2.364^{[-4]})$	0.059 $(0.016)^*$	0.926 $(0.018)^*$	59093.055

Entries of the form $x^{[-y]}$ represent scientific notation; that is $2.5^{[-2]}$ represents the value 2.5×10^{-2}.

The table presents the parameter estimates for the nested GARCH models constructed to test for model misspecification in the functional form of the conditional mean equation. The first row presents the parameter estimated for the model which imposes no restrictions on the augmented conditional mean equation (Equation 4.8 in the text). The following five rows present the parameter estimate for the models with all reasonable subsets of conditional mean equation parameters included. Numbers in parentheses under the parameter estimates represent asymptotic standard errors. Parameters which are significantly different from zero at the 5% level are indicated with an asterisk. The data set comprises excess returns on an U.S. equity index covering the period January 2, 1926 to December 29, 1990, a total of 17,484 observations.

Table D.2: Descriptive statistics for standardized residuals (alternative mean equations)

Model	Mean	Variance	Skewness	Kurtosis	Persistence	Sign Bias	Negative	Positive
model25	-1.508×10^{-4}	1.000	−0.533	4.913	0.96247	1.411	−2.941	−0.533
model18	-1.174×10^{-3}	1.000	−0.567	4.913	0.96361	2.567	−3.555	−1.403
model27	1.419×10^{-3}	1.000	−0.554	5.083	0.96033	2.080	−2.825	−0.904
model28	3.207×10^{-4}	1.000	−0.560	5.069	0.96125	2.237	−2.986	−0.999
model26	9.112×10^{-4}	1.000	−0.527	4.913	0.96164	1.003	−2.785	−0.445
model22	-1.595×10^{-4}	1.000	−0.533	4.913	0.96247	1.411	−2.941	−0.533

The table presents additional descriptive statistics for the standardized residuals from the mean equation for the models presented in Table D.1. The first through fourth columns present the first four statistical moments of the standardized residuals from the mean equation. The fifth column presents the Jarque and Bera (1987) test for normality. The sixth column presents the first order autocorrelation of the square root of the conditional variance of the standardized residuals, which is a measure of the persistence of the conditional volatility. The seventh through ninth columns present the results of the sign-bias tests suggested by Engle and Ng (1993), explained in Section 3.3 of the text. The data set comprises excess returns on an U.S. equity index covering the period January 2, 1926 to December 29, 1990, a total of 17,484 observations.

Table D.3: Estimates of model parameters (with regimes)

Model	γ_0	γ_1	λ	ν	b	c	ω	α	β	Likelihood
model01	$4.929^{[-4]}$ $(8.734^{[-5]})^*$	1.677 (1.227)	0	1	0	0	−0.247 (0.030)*	0.066 (0.008)*	0.960 (0.006)*	58815.024
			0	0	0	0	−0.063 (0.013)*	0.059 (0.016)*	0	
			0	0	0	0	0	0.082 (0.017)*	0.007 (0.002)*	
model02	$5.582^{[-4]}$ $(6.989^{[-5]})^*$	−2.212 (0.908)*	0	1	0.453 (0.053)*	0	−0.424 (0.074)*	0.085 (0.011)*	0.928 (0.014)*	59059.888
			0	0	0.297 (0.058)*	0	0	0.042 (0.011)*	0.016 (0.003)*	
			0	0	0.222 (0.076)*	0	0.162 (0.082)	0.053 (0.019)*	0.041 (0.015)*	
model03	$4.926^{[-4]}$ $(7.544^{[-5]})^*$	−0.943 (1.189)	0	1	0	0.514 (0.057)*	−0.346 (0.071)*	0.061 (0.009)*	0.939 (0.014)*	59044.833
			0	0	0	0.219 (0.112)	0	0.025 (0.011)*	0.009 (0.002)*	
			0	0	0	0	0.149 (0.079)	0.050 (0.017)*	0.034 (0.015)*	
model04	$5.179^{[-4]}$ $(6.788^{[-5]})^*$	−1.862 (0.884)*	0	1	−0.458 (0.025)*	0.814 (0.098)*	−0.256 (0.029)*	0.056 (0.006)*	0.953 (0.006)*	59079.940
			0	0	1.232 (0.036)*	−0.870 (0.130)*	−0.113 (0.015)*	0.070 (0.018)*	0	
			0	0	0.899 (0.059)*	−0.600 (0.121)*	0	0.070 (0.013)*	0.014 (0.002)*	

Table D.3: continued.

Model	γ_0	γ_1	λ	ν	b	c	ω	α	β	Likelihood
model05	$5.182^{[-4]}$ $(7.716^{[-5]})^*$	1.399 (0.856)	1	1	0	0	$3.628^{[-4]}$ $(5.981^{[-5]})^*$	0.073 (0.010)*	0.898 (0.013)*	58798.933
			0	0	0	0	0	0.080 (0.018)*	−0.085 (0.015)*	
			0	0	0	0	0	0.074 (0.022)*	−0.028 (0.015)	
model06	$5.625^{[-4]}$ $(6.403^{[-5]})^*$	−2.219 (0.788)*	1	1	0.443 (0.045)*	0	$4.844^{[-4]}$ $(7.639^{[-5]})^*$	0.084 (0.011)*	0.867 (0.016)*	59056.272
			0	0	0.323 (0.053)*	0	$-1.702^{[-4]}$ $(9.537^{[-5]})$	0.053 (0.018)*	−0.066 (0.029)*	
			0	0	0.223 (0.058)*	0	0	0.064 (0.017)*	−0.034 (0.014)*	
model07	$4.980^{[-4]}$ $(6.927^{[-5]})^*$	−0.963 (1.149)	1	1	0	0.548 (0.055)*	$5.167^{[-4]}$ $(8.155^{[-5]})^*$	0.071 (0.007)*	0.879 (0.013)*	59046.200
			0	0	0	0.239 (0.107)*	$-2.490^{[-4]}$ $(5.855^{[-5]})^*$	0.017 (0.009)	0	
			0	0	0	0	0	0.049 (0.008)*	0	
model08	$5.501^{[-4]}$ $(9.518^{[-5]})^*$	−2.099 (1.478)	1	1	−0.123 (0.022)*	0.629 (0.093)*	$4.753^{[-4]}$ $(8.850^{[-5]})^*$	0.065 (0.009)*	0.893 (0.016)*	59078.830
			0	0	0.915 (0.029)*	−0.697 (0.107)*	$-1.652^{[-4]}$ $(1.043^{[-4]})$	0.080 (0.018)*	−0.094 (0.027)*	
			0	0	0.562 (0.036)*	−0.374 (0.113)*	0	0.065 (0.015)*	−0.044 (0.012)*	

Table D.3: continued.

Model	γ_0	γ_1	λ	ν	b	c	ω	α	β	Likelihood
model09	$5.204^{[-4]}$ $(6.884^{[-5]})^*$	1.239 (0.739)	2	2	0	0	$2.229^{[-6]}$ $(4.385^{[-7]})^*$	0.054 (0.009)*	0.909 (0.014)*	58847.343
			0	0	0	0	0	0.099 (0.019)*	−0.140 (0.022)*	
			0	0	0	0	$3.827^{[-6]}$ $(1.429^{[-6]})^*$	0.089 (0.025)*	−0.061 (0.023)*	
model10	$6.059^{[-4]}$ $(6.725^{[-5]})^*$	−2.929 (0.844)*	2	2	0.758 (0.057)*	0	$4.854^{[-6]}$ $(8.592^{[-7]})^*$	0.060 (0.007)*	0.823 (0.020)*	59090.123
			0	0	0	0	$-3.230^{[-6]}$ $(9.124^{[-7]})^*$	0.050 (0.015)*	−0.054 (0.035)	
			0	0	0	0	0	0.050 (0.009)*	0	
model11	$5.514^{[-4]}$ $(6.654^{[-5]})^*$	−1.643 (0.796)*	2	2	0	0.542 (0.087)*	$4.536^{[-6]}$ $(6.526^{[-7]})^*$	0.043 (0.007)*	0.865 (0.013)*	59024.161
			0	0	0	0	$-3.375^{[-6]}$ $(5.377^{[-7]})^*$	0.025 (0.007)*	0	
			0	0	0	−0.206 (0.099)*	0	0.075 (0.011)*	0	
model12	$6.062^{[-4]}$ $(6.715^{[-5]})^*$	−2.920 (0.842)*	2	2	0.819 (0.092)*	−0.053 (0.060)	$4.928^{[-6]}$ $(8.661^{[-7]})^*$	0.065 (0.009)*	0.818 (0.021)*	59090.727
			0	0	0	0	$-3.304^{[-6]}$ $(9.174^{[-7]})^*$	0.053 (0.017)*	−0.054 (0.036)	
			0	0	0	0	0	0.053 (0.010)*	0	

Table D.3: continued.

Model	γ_0	γ_1	λ	ν	b	c	ω	α	β	Likelihood
model13	$5.297^{[-4]}$ $(7.022^{[-5]})^*$	1.272 (0.768)	1.774 $(0.818)^*$	1.774	0	0	$8.007^{[-6]}$ $(3.277^{[-5]})$	0.063 (0.038)	0.899 $(0.029)^*$	58849.615
			0	0	0	0	$-2.582^{[-6]}$ $(9.660^{[-6]})$	0.083 (0.049)	−0.091 (0.054)	
			0	0	0	0	$7.975^{[-6]}$ $(2.612^{[-5]})$	0.087 $(1.687^{[-5]})^*$	−0.047 $(2.681^{[-7]})^*$	
model14	$5.605^{[-4]}$ $(6.533^{[-5]})^*$	−2.067 $(0.800)^*$	1.941 $(0.127)^*$	1.941	0.740 $(0.052)^*$	0	$6.965^{[-6]}$ $(3.809^{[-6]})$	0.067 $(0.009)^*$	0.813 $(0.022)^*$	59092.546
			−0.260 $(0.028)^*$	−0.260	0	0	0	0.040 $(0.008)^*$	0	
			0	0	0	0	0	0.054 $(0.010)^*$	0	
model15	$5.365^{[-4]}$ $(6.855^{[-5]})^*$	−1.513 (0.939)	1.471 (2.368)	1.471	0	0.447 $(0.138)^*$	$5.099^{[-5]}$ $(6.219^{[-4]})$	0.063 (0.106)	0.878 $(0.045)^*$	59052.350
			0	0	0	0.465 (0.534)	$-2.950^{[-5]}$ $(3.303^{[-4]})$	0	0	
			0	0	0	0	0	0.056 $(3.905^{[-8]})^*$	0	
model16	$5.621^{[-4]}$ $(6.561^{[-5]})^*$	−1.988 $(0.803)^*$	1.900 $(0.119)^*$	1.900	0.707 $(0.109)^*$	0.012 (0.078)	$8.530^{[-6]}$ $(4.279^{[-6]})$	0.068 $(0.012)^*$	0.814 $(0.024)^*$	59092.840
			−0.261 $(0.028)^*$	−0.261	0	0	0	0.043 $(0.009)^*$	0	
			0	0	0	0	0	0.056 $(0.011)^*$	0	

Table D.3: continued.

Model	γ_0	γ_1	λ	ν	b	c	ω	α	β	Likelihood
model17	$5.301^{[-4]}$	1.223	0.839	1.677	0	0	$8.687^{[-4]}$	0.034	0.919	58846.266
	$(7.016^{[-5]})$*	(0.777)	(0.294)*	(0.115)*			(0.001)	(0.012)*	(0.015)*	
			0	0	0	0	0	0.037	−0.054	
								(0.018)*	(0.025)*	
			0	0	0	0	0	0.030	0	
								(0.010)*		
model18	$6.030^{[-4]}$	−2.761	5.879	2.683	0.754	0	$2.167^{[-14]}$	0.112	0.703	59108.460
	$(7.029^{[-5]})$*	(0.896)*	(1.820)*	(0.390)*	(0.087)*		$(1.943^{[-13]})$	(0.033)*	(0.070)*	
			−4.068	−1.548	0	0	$4.385^{[-6]}$	0.117	0	
			(1.578)*	(0.447)*			$(1.331^{[-5]})$	(0.057)*		
			−4.332	−0.989	0	0	$3.967^{[-5]}$	0	0.134	
			(1.754)*	(0.477)*			$(1.102^{[-4]})$		(0.054)*	
model19	$5.607^{[-4]}$	−1.883	4.116	1.533	0	0.472	$1.013^{[-10]}$	0.182	0.792	59066.704
	$(7.537^{[-5]})$*	(0.945)	(2.009)*	(0.184)*		(0.064)*	$(9.964^{[-10]})$	(0.089)*	(0.077)*	
			−2.510	−0.613	0	0.322	$1.473^{[-5]}$	0	0	
			(1.344)	(0.228)*		(0.104)*	$(7.373^{[-5]})$			
			−1.866	0	0	0	$1.578^{[-6]}$	0	0.042	
			(1.072)				$(7.803^{[-6]})$		(0.019)*	
model20	$5.977^{[-4]}$	−2.701	5.684	2.659	0.771	−0.024	$5.680^{[-14]}$	0.114	0.707	59108.625
	$(6.609^{[-5]})$*	(0.848)*	(2.021)*	(0.407)*	(0.032)*	(0.043)	$(5.665^{[-13]})$	(0.034)*	(0.071)*	
			−3.927	−1.536	0	0	$5.793^{[-6]}$	0.113	0	
			(1.704)*	(0.436)*			$(1.787^{[-5]})$	(0.055)*		
			−4.115	−0.945	0	0	$3.627^{[-5]}$	0	0.128	
			(1.910)*	(0.496)			$(1.051^{[-4]})$		(0.053)*	

Entries of the form $x^{[-y]}$ represent scientific notation; that is $2.5^{[-2]}$ represents the value 2.5×10^{-2}.

The table presents the parameter estimates for the regime dependent versions of the nested GARCH models which were presented in conventional form in Table 3.2. Numbers in parentheses under the parameter estimates represent asymptotic standard errors. Parameters which are significantly different from zero at the 5% level are indicated with an asterisk. The first pair of rows for each model present the parameter estimates and standard errors for the base regime, and the remaining four rows present the parameter estimates and standard errors for the regimes of unusually low and high volatility respectively. Parameters which are excluded by the stepwise estimation procedure are represented by a value of zero, with no corresponding standard error presented. The data set comprises excess returns on an U.S. equity index covering the period January 2, 1926 to December 29, 1990, a total of 17,484 observations.

Table D.4: Descriptive statistics for standardized residuals (with regimes)

Model	Mean	Variance	Skewness	Kurtosis	Persistence	Sign Bias	Negative	Positive
model01	-4.185×10^{-1}	0.998	−0.492	3.217	0.97618	7.436	−6.621	−5.079
	-3.376×10^{-2}	1.001	−0.568	4.042	0.97618			
	-2.669×10^{-2}	0.989	−0.410	1.069	0.95831			
	-7.167×10^{-2}	0.999	−0.382	2.947	0.96390			
model02	-4.783×10^{-2}	1.000	−0.408	3.318	0.94195	1.799	−2.808	−0.899
	-7.626×10^{-3}	1.000	−0.546	4.536	0.94195			
	4.959×10^{-3}	0.998	−0.212	0.783	0.95610			
	-5.686×10^{-3}	1.003	−0.243	2.456	0.96276			
model03	-8.650×10^{-2}	1.000	−0.429	3.385	0.95109	3.098	−2.523	−1.000
	-8.661×10^{-3}	0.999	−0.567	4.736	0.95109			
	8.859×10^{-3}	0.997	−0.253	0.763	0.95795			
	-2.184×10^{-2}	1.003	−0.250	2.294	0.96628			
model04	-4.187×10^{-2}	1.000	−0.418	3.161	0.96475	3.123	−2.922	−0.912
	-6.735×10^{-3}	0.999	−0.558	4.328	0.96475			
	9.270×10^{-3}	0.999	−0.231	0.776	0.96182			
	-8.562×10^{-3}	1.003	−0.241	2.308	0.96100			
model05	-4.240×10^{-1}	0.998	−0.500	3.281	0.97351	7.286	−6.968	−5.093
	-3.478×10^{-2}	0.999	−0.559	3.993	0.97351			
	-2.962×10^{-2}	0.992	−0.382	1.026	0.94473			
	-6.934×10^{-2}	1.002	−0.454	3.347	0.97563			
model06	-5.102×10^{-2}	1.000	−0.411	3.356	0.95553	1.684	−3.481	−0.899
	-8.504×10^{-3}	1.000	−0.549	4.523	0.95553			
	6.325×10^{-3}	0.999	−0.187	0.783	0.95870			
	-6.007×10^{-3}	1.001	−0.265	2.630	0.96672			
model07	-8.453×10^{-2}	1.000	−0.421	3.411	0.94922	2.707	−2.563	−0.707
	-8.523×10^{-3}	1.004	−0.557	4.721	0.94922			
	9.229×10^{-3}	0.995	−0.240	0.767	0.96608			
	-2.164×10^{-2}	0.995	−0.247	2.377	0.96923			
model08	-4.655×10^{-2}	1.000	−0.419	3.384	0.95746	2.169	−2.803	−0.679
	-7.622×10^{-3}	0.999	−0.561	4.661	0.95746			
	7.104×10^{-3}	0.999	−0.194	0.777	0.95865			
	-6.798×10^{-3}	1.002	−0.267	2.445	0.96719			

Table D.4: continued.

Model	Mean	Variance	Skewness	Kurtosis	Persistence	Sign Bias	Negative	Positive
model09	-4.085×10^{-1}	0.998	−0.471	3.141	0.97782	7.339	−5.912	−4.953
	-3.399×10^{-2}	1.000	−0.573	4.045	0.97782			
	-2.819×10^{-2}	0.994	−0.361	0.967	0.93344			
	-6.597×10^{-2}	0.998	−0.321	2.727	0.96821			
model10	-3.268×10^{-2}	1.000	−0.397	3.245	0.93763	2.003	−1.618	−0.739
	-6.928×10^{-3}	0.999	−0.536	4.492	0.93763			
	-1.376×10^{-3}	0.995	−0.218	0.794	0.95341			
	3.610×10^{-3}	1.007	−0.220	2.252	0.96848			
model11	-1.063×10^{-1}	1.000	−0.407	3.054	0.94031	4.164	−2.152	−1.753
	-1.105×10^{-2}	0.998	−0.536	4.217	0.94031			
	-4.349×10^{-3}	0.995	−0.293	0.813	0.96362			
	-1.439×10^{-2}	1.008	−0.204	2.105	0.96780			
model12	-3.324×10^{-2}	1.000	−0.397	3.258	0.93653	1.925	−1.656	−0.764
	-6.918×10^{-3}	0.999	−0.535	4.504	0.93653			
	-1.281×10^{-3}	0.996	−0.215	0.795	0.95332			
	3.289×10^{-3}	1.007	−0.221	2.278	0.96828			
model13	-4.240×10^{-1}	0.998	−0.474	3.125	0.97399	7.353	−6.001	−5.046
	-3.534×10^{-2}	0.999	−0.565	3.986	0.97399			
	-3.073×10^{-2}	0.996	−0.379	0.985	0.94666			
	-6.725×10^{-2}	0.998	−0.340	2.792	0.96935			
model14	-5.561×10^{-2}	0.999	−0.396	3.229	0.93219	1.867	−1.358	−0.817
	-7.067×10^{-3}	0.998	−0.532	4.483	0.93219			
	1.261×10^{-3}	0.994	−0.237	0.777	0.96494			
	-7.290×10^{-3}	1.006	−0.210	2.224	0.96603			
model15	-9.132×10^{-2}	0.999	−0.415	3.207	0.94918	3.571	−2.309	−1.311
	-1.047×10^{-2}	0.999	−0.554	4.458	0.94918			
	2.672×10^{-3}	0.997	−0.251	0.766	0.96532			
	-1.499×10^{-2}	1.002	−0.223	2.203	0.97018			
model16	-6.568×10^{-2}	0.999	−0.397	3.225	0.93206	1.869	−1.341	−0.886
	-7.950×10^{-3}	0.998	−0.532	4.474	0.93206			
	4.965×10^{-4}	0.994	−0.238	0.782	0.96460			
	-8.765×10^{-3}	1.004	−0.210	2.220	0.96567			

Table D.4: continued.

Model	Mean	Variance	Skewness	Kurtosis	Persistence	Sign Bias	Negative	Positive
model17	-4.191×10^{-1}	0.998	−0.472	3.153	0.96051	7.307	−6.045	−4.861
	-3.491×10^{-2}	1.000	−0.561	3.975	0.96051			
	-3.048×10^{-2}	0.996	−0.374	0.990	0.93285			
	-6.644×10^{-2}	0.996	−0.345	2.921	0.96597			
model18	-4.449×10^{-2}	1.000	−0.390	3.170	0.95759	1.890	−2.166	−0.631
	-8.527×10^{-3}	0.999	−0.536	4.333	0.95759			
	9.703×10^{-4}	0.996	−0.179	0.786	0.95949			
	7.162×10^{-4}	1.006	−0.220	2.313	0.96709			
model19	-7.980×10^{-2}	0.999	−0.409	3.214	0.96444	3.125	−2.632	−1.225
	-1.106×10^{-2}	0.996	−0.556	4.465	0.96444			
	4.176×10^{-3}	1.004	−0.205	0.773	0.95806			
	-1.016×10^{-2}	1.003	−0.234	2.241	0.97349			
model20	-4.355×10^{-2}	1.000	−0.391	3.175	0.95700	1.884	−2.172	−0.703
	-8.267×10^{-3}	0.999	−0.536	4.343	0.95700			
	1.604×10^{-3}	0.998	−0.180	0.786	0.95951			
	2.463×10^{-5}	1.005	−0.221	2.318	0.96701			

The table presents additional descriptive statistics for the standardized residuals from the mean equation for the models presented in Table D.1. For each model presented in the table, The first row presents the statistics for the entire set of standardized residuals. The second row of this section presents the statistics for the residuals from the base regime only, and the final two rows present the statistics for the low and high volatility regimes respectively. In all rows of the table, the first through fourth columns present the first four statistical moments of the standardized residuals from the mean equation. The fifth column presents the Jarque and Bera (1987) test for normality. The sixth column presents the first order autocorrelation of the square root of the conditional variance of the standardized residuals, which is a measure of the persistence of the conditional volatility. The seventh through ninth columns present the results of the sign-bias tests suggested by Engle and Ng (1993) and explained in Section 3.3 of the text. The sign-bias test results are not reported for the final three rows of each model as the segmented nature of the regime results makes the application of these tests impossible. The data set comprises excess returns on an U.S. equity index covering the period January 2, 1926 to December 29, 1990, a total of 17,484 observations.